

iBARLE: imBalance-Aware Room Layout Estimation

Taotao Jing¹, Lichen Wang², Naji Khosravan², Zhiqiang Wan², Zachary Bessinger²,
Zhengming Ding¹, Sing Bing Kang²

¹Tulane University ²Zillow Group

{tjing, zding1}@tulane.edu

{lichenw, najik, zhiqiangw, zacharybe, singbingk}@zillowgroup.com

Abstract

Room layout estimation predicts layouts from a single panorama. It requires datasets with large-scale and diverse room shapes to well train the models. However, there are significant imbalances in real-world datasets including the dimensions of layout complexity, camera locations, and variation in scene appearance. These issues considerably influence the model training performance. In this work, we propose *imBalance-Aware Room Layout Estimation (iBARLE)* framework to address these issues. *iBARLE* consists of: (1) *Appearance Variation Generation (AVG)* module, which promotes visual appearance domain generalization, (2) *Complex Structure Mix-up (CSMix)* module, which enhances generalizability w.r.t. room structure, and (3) a gradient-based layout objective function, which allows more effective accounting for occlusions in complex layouts. All modules are jointly trained and help each other to achieve the best performance. Experiments and ablation studies based on *ZInD* [6] dataset illustrate that *iBARLE* has state-of-the-art performance compared with other layout estimation baselines.

1. Introduction

With the recent advancements in computer vision related applications (e.g., AR/VR, virtual touring, and navigation), room layout estimation is receiving a lot of attention from researchers. Specifically, panorama-based layout estimation has been a major area of focus due to the increased 360° field of view [49]. A great deal of progress was made on monocular layout estimation based on a single panorama [38, 48, 39]. Some directly use the equirectangular panorama [30] while others combine the equirectangular panorama with its perspective top-down view [42]. A recent trend of papers formulate the problem as a 1-dimension sequence that represents depth on the horizon line of the panorama and calculates the room height by the consistency between the horizon-depth of ceiling and floor boundaries




Figure 1. Imbalanced and diverse sample distributions such as various room corner numbers, appearances, Manhattan styles, and capture locations which downgrade estimation performances.

[41]. Some methods directly predict the room height to make better geometry awareness of the room layout in the vertical direction [19].

However, these techniques are less effective for complex room shapes (e.g., self-occlusion and non-Manhattan) [42]. As a result, the majority of these approaches conform to the Manhattan World or Atlanta World assumption, as well as their corresponding post-processing strategies, resulting in promising performances for simple rooms, while failing to achieve the same level of performance for complex layouts. For instance, *ZInD* [6] is a very large-scale indoor dataset. The performance of most single-panorama layout estimation solutions degrades in complex room shape scenarios. On top of that, there are implicit data imbalance challenges present in real-world datasets such as different camera poses, illumination changes, and texture variations (see Figure 1). As shown in Figure 4, simple rooms with only four corners take the majority of the dataset while complex rooms with nine corners make up only 2% of the dataset. Consequently, models trained on such datasets will have a tendency to get biased towards the majority.

In this work, we address the data imbalance and appearance variation issues associated with room layout estimation. We specifically design an *Appearance Variation Generalization (AVG)* module to overcome appearance variations that are present in real-world datasets (e.g., pose, tex-

ture, illumination), a Complex Structure Mix-up (CSMix) module to handle the long tail imbalanced distribution of structural complexity of the training and test data, and a layout gradient-based cost function to better deal with occlusions in complex indoor spaces. The contributions of our work are listed below:

- To the best of our knowledge, our work is the first to tackle the data imbalance issue in layout estimation systematically and based on different sources of bias in the indoor dataset.
- Our framework, iBARLE, introduces architectural and training novelties through our proposed Appearance Variation Generalization (AVG) module, Complex Structure Mix-up (CSMix) module, and our layout gradient-based cost function.
- We show iBARLE consistently achieves the state-of-the-art performance for both simple/conventional and complex real-world indoor layout estimation tasks through experiments and ablation studies.

2. Related Work

Indoor Panoramic Layout Estimation. Most panorama-based indoor layout estimation efforts use simple indoor room shapes with Manhattan World [5] assumptions and Atlanta World [34] assumptions and corresponding pose-processing operations. Specifically, convolutional neural networks (CNNs) take panorama images as input to extract visual features which will be used to estimate layout. With LayoutNet, layout estimation of a cuboid room with Manhattan constraints is reconstructed based on images aligned with vanishing points and detected layout elements such as corners and boundaries, and the model designed for cuboid rooms is further extended to predicting Manhattan layouts in general [51]. In contrast, Dula-Net projects the panorama images into two different views, equirectangular and perspective, to predict the floor and ceiling probability maps and two-dimensional floor plans [48].

A new dataset containing panoramas of Manhattan-world room layouts with different numbers of corners is introduced as Realtor360 to learn more complex room layouts [48]. Moreover, HorizonNet outperforms other strategies by representing a 2D room layout as three 1-D vectors at each image column, and the 1-D sequence vectors encode the positions of floor-wall and ceiling-wall, and the existence of wall-wall boundaries [38]. Although the two improved frameworks LayoutNet v2 and Dula-Net v2 achieve better performance on cuboid datasets, HorizonNet has been demonstrated to be more effective on the new MatterportLayout dataset for general Manhattan layout estimation tasks [52]. AtlantaNet breaks through the Manhattan World limitations and projects the original gravity-aligned panorama images on two horizontal planes to reconstruct the Atlanta World 3D bounding surfaces of the rooms [30].

HoHoNet is the first work exploring compact latent horizontal features learning for efficient and accurate layout reconstruction and depth estimation equipped with Efficient Height Compression (EHC) and multi-head self-attention (MSA) modules [39]. PSMNet is a pioneering end-to-end joint layout-pose deep architecture for large and complex room layout estimation from a pair of panoramas [42]. LED²-Net goes beyond conventional 3D layout estimation by predicting depth on the horizon line of the panorama, and a differentiating depth rendering procedure is proposed to maximize the 3D geometric information without needing to provide the ground-truth depth [41]. LGT-Net further extends the LED²-Net with an SWG-Transformer module, which consists of shifted window blocks and global blocks, to predict both horizon depth and room height with a planar geometry aware loss to supervise the estimation of the planeness of walls and turning of corners [19].

Imbalanced Data Training. Due to the widespread application of artificial intelligence systems in our daily lives, the issue of data imbalance has gained increasing attention in recent years [27]. Research about imbalance AI seeks to ensure that AI systems make decisions and predictions without discriminating against certain groups of data when making crucial and life-changing decisions. The algorithms that target bias generally fall into three categories: (1) Pre-processing methods attempt to remove the underlying discrimination from the data [1, 7]. (2) In-processing techniques seek to modify the learning and training strategies, either by incorporating changes into the objective function or imposing a constraint, to eliminate discrimination during model training [1, 4]. (3) Post-processing is performed after training stage if the algorithm treats the AI model as a black box without modifying the training data or learning strategies. Specifically, a hold-out set that was not directly involved in the training process is used during the post-processing phase to reassign the initial labels assigned by the black-box model [1, 4].

Domain Generalization. As part of the quest to develop models that can generalize to unknown distributions, Domain Generalization (DG), which refers to out-of-distribution generalization, has attracted increasing attention in recent years [43]. A majority of existing domain generalization methods can be classified into three types: (1) Data manipulation focuses on manipulating the inputs in order to assist in the learning of general representations. Data augmentation [36, 15] and data generalization [31, 24, 32] are two types of popular techniques in this regard. Additionally, mix-up on the original image-level [44, 45] or feature-level [50, 47] is a popular and effective technique. (2) Representation learning [3] consists of two representative techniques: a). Learning domain-invariant representations

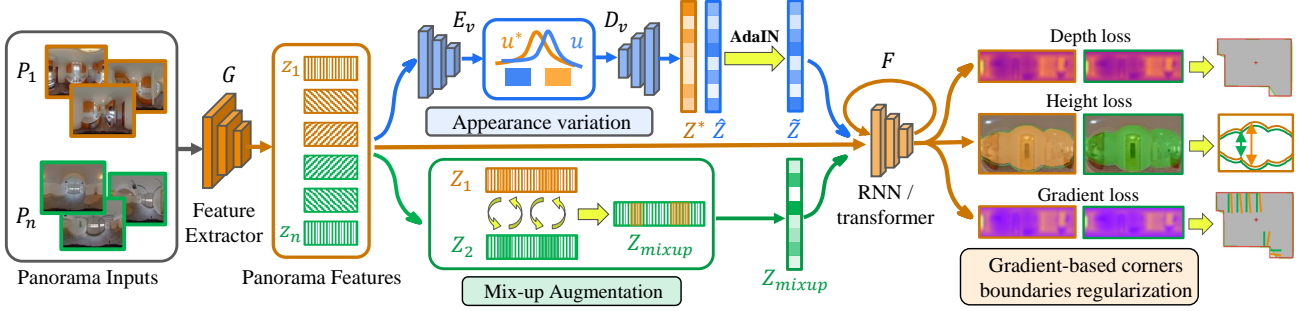


Figure 2. Framework of iBARLE. There are four core modules: visual feature extractor $G(\cdot)$, transformer-based sequential layout estimation $F(\cdot)$, Appearance Variation Generalization (AVG), and Complex Structure Mix-up (CSMix). AVG enhances the generalization capability by enhancing invariance to appearance changes. CSMix improves the prediction of complex room shapes by generating more diverse layouts that improves the balance of the overall sample distributions.

via kernel functions [13, 2], adversarial learning [11, 12], feature alignment [28, 20, 23], etc.; b). Feature disentanglement is the process of distancing features into domain-shared and/or domain-specific elements to enhance generalization [21, 29, 8]. (3) Learning strategies that promote generalization employ general strategies such as ensemble learning [9, 46, 26], meta-learning [10, 37, 33], gradient operations [35, 17], and self-supervised learning [22, 18, 25]. To the best of our knowledge, we are the first to apply the domain generalization techniques to address the appearance variation issues associated with room layout estimation.

3. Framework Overview

The iBARLE framework, illustrated in Figure 2, has four core modules: visual feature extractor $G(\cdot)$, transformer-based sequential layout estimation module $F(\cdot)$, Appearance Variation Generalization module (AVG), and Complex Structure Mix-up module (CSMix). The room layout is estimated by predicting the horizon depths and heights of multiple points sampled from the floor boundary of the polygon of the panorama image [41, 19]. During the training phase, the AVG and CSMix modules are used to compensate for the imbalance in the data; during the test phase, only the feature extraction and the sequential layout estimation module are called to predict the room layout.

4. Proposed Algorithm

In this section, we first go over the overall framework to predict the room layout based on a single panorama input. We then introduce the two modules, Appearance Variation Generalization (AVG) and Complex Structure Mix-up (CSMix). Finally, the gradient-based *Corners and Occlusions objective* is presented. It improves the efficiency of the model in detecting visible wall-wall corners and occlusion boundaries in complex spatial arrangements.

4.1. Framework

Reconstructing the 3D indoor layout by predicting the horizon depth and room height is an effective strategy that is widely adopted by recent methods [19, 41]. Specifically, N points $\mathbf{P} = \{p_i\}_{i=1}^N$ with equal longitude interval are sampled from the floor-boundary of the polygon of the panorama image, where the longitudes of the sampled points are $\theta_i = 2\pi(\frac{i}{N} - 0.5)g_{i=1}^N$. Then, the coordinate of the point p_i and the corresponding horizon-depth d_i can be obtained as:

$$\begin{aligned} \mathbf{p}_i &= (x_i, y_i, z_i), \\ d_i &= \sqrt{x_i^2 + y_i^2}. \end{aligned} \quad (1)$$

The sampled points $\mathbf{P} = \{p_i\}_{i=1}^N$ can be converted into horizon-depth sequence $\{d_i\}_{i=1}^N$. In addition, the height h of each room is adopted to further supervise the model prediction on the vertical direction.

As illustrated in Figure 2, the visual feature sequence extracted by the feature extractor $G(\cdot)$ is passed to a sequential neural networks $F(\cdot)$ to predict the horizon-depth $\hat{\mathbf{d}} = \{\hat{d}_i\}_{i=1}^N$ and room height sequence $\hat{\mathbf{h}} = \{\hat{h}_i\}_{i=1}^N$, respectively, through two separate branches in the output layer. A limitation of prior work is the assumption that the room height is always the same in the indoor space. Differently, we predict \hat{h}_i for each sampled point $p_i \in \mathbf{P}$ with ground-truth height denoted as $h_i \in \mathbf{h}$, where $\mathbf{h} = \{h_i\}_{i=1}^N$, which fits the design of our proposed structure variation generalization module (Introduced in Section 4.3). Our sequential neural networks is a SWG-Transformer module following [19]. The predicted horizon-depth and room height are supervised by ground-truth formulated as:

$$\begin{aligned} L_d &= \frac{1}{N} \sum_{i=1}^N j d_i - \hat{d}_i j, \\ L_h &= \frac{1}{N} \sum_{i=1}^N j h_i - \hat{h}_i j, \end{aligned} \quad (2)$$




Figure 3. Corners and occlusions awareness constraint.

where \hat{d}_i and d_i are the predicted and ground-truth horizon-depth at point p_i , respectively, and \hat{h}_i and h_i are the room height prediction and ground-truth, respectively.

4.2. Appearance Variation Generalization (AVG)

Capturing panoramas in different situations (e.g., illuminations, wall/carpet colors, cameras, and textures) would significantly change the visual appearances of the panorama images. Inspired by domain generalization problems and techniques, we propose AVG to disentangle appearance variations of panoramas during the training. This encourages the model to focus only on the important indoor structural information for the downstream task and ignore the rest. Subsequently, the model provides a better generalization capacity to novel appearances.

Specifically, the visual features \mathbf{Z} extracted from $G(\cdot)$ are input to an appearance encoder $E_v(\cdot)$ which projects the visual features into a Normal distribution $\mathbf{u} \sim \mathcal{N}(0, \mathbf{I})$, then a decoder $D_v(\cdot)$ project the latent embeddings \mathbf{u} into the original visual features space as $\hat{\mathbf{Z}}$ by minimizing a regularized norm-based loss:

$$L_{var} = k\mu(\mathbf{Z}) - \mu(\hat{\mathbf{Z}})k_2 + k\sigma(\mathbf{Z}) - \sigma(\hat{\mathbf{Z}})k_2. \quad (3)$$

To further enhance the generalizability of the model to novel appearance, we randomly sample latent embeddings $\mathbf{u}^* \sim \mathcal{N}(0, \mathbf{I})$ and project them into the visual features space as \mathbf{Z}^* , which is a novel appearance variation never observed in the training data. Then, the style information of the randomly sampled feature \mathbf{Z}^* is transferred to the real content sample \mathbf{Z} via AdaIN [16] strategy as ([50] similar to style mixup, [25] Feature Stylization):

$$\tilde{\mathbf{Z}} = \sigma(\mathbf{Z}^*) \left(\frac{\mathbf{Z} - \mu(\mathbf{Z})}{\sigma(\mathbf{Z})} \right) + \mu(\mathbf{Z}^*), \quad (4)$$

where the synthetic sample $\tilde{\mathbf{Z}}$ is then input to the following sequential module to predict the horizon depth and room height and optimize the model as Eq. (2) since they share the same content and spatial structure information and layout.

4.3. Complex Structure Mix-up (CSMix)

Another main issue that exists in prior works is over-fitting to simple room shapes (e.g., Manhattan world or




Figure 4. Statistics of ZInD dataset split by different attributes.

Atlanta world pre-assumption), and cannot handle minority and complex room shapes. Thus, in order to enhance the training data space complexity, some novel and complex samples are synthesized via cross-room, column-based mix-up strategy. Specifically, for two randomly paired samples $\tilde{\mathbf{Z}}_i, \mathbf{Z}_j \mathcal{G}$, we randomly exchange columns across these two samples, which will result in two new samples with more complex structures. It is noteworthy that the synthesized samples contain more complex layout structures and also mixed appearance variations due to the cross-room mix-up. However, to avoid the random mix-up producing unreasonable noisy samples harming the optimization, a sequence of consecutive columns with width $1 \leq w \leq N$ is randomly selected from the two samples and mix-up, where N is the width, i.e., the number of columns, of the whole visual feature maps:

$$\begin{aligned} \tilde{\mathbf{Z}}_i &= \mathbf{Z}_i[:c_i] \quad \mathbf{Z}_j[c_j:c_j+w] \quad \mathbf{Z}_i[c_i+w:], \\ \tilde{\mathbf{Z}}_j &= \mathbf{Z}_j[:c_j] \quad \mathbf{Z}_i[c_i:c_i+w] \quad \mathbf{Z}_j[c_j+w:], \end{aligned} \quad (5)$$

where (\cdot, \cdot) is concatenation operation, and $1 \leq c_i \leq N - w$ and $1 \leq c_j \leq N - w$ are randomly sampled points or columns start indices for \mathbf{Z}_i and \mathbf{Z}_j , respectively. Similarly, the synthesized samples $\tilde{\mathbf{Z}}_{i/j}$ are also input to the following layout estimation to improve the generalizability of the whole framework. It is noteworthy that since the mix-up synthesis is applied across rooms, so the room heights of the synthesized samples are also mixed up, e.g., $\tilde{\mathbf{h}}_i = \mathbf{h}_i[:c_i] \quad \mathbf{h}_j[c_j:c_j+w] \quad \mathbf{h}_i[c_i+w:]$. Then the model is optimized through learning objectives as Eq. (2).

4.4. Corners and Occlusions Awareness Constraint

As shown in Figure 3, each wall is a plane but the positions on the same wall could have different horizon depths. Thus, to supervise the planar of the walls, the normals at different positions of the same wall are constrained consistently [19]. Since the wall and floors are preassumed perpendicular to the floor and the normals. Thus we first convert the ground-truth and predicted horizon depth at each position back into the 3D point $\mathbf{p}_i = (d_i \sin(\theta_i), h^f, d_i \cos(\theta_i))$ and $\hat{\mathbf{p}}_i = (\hat{d}_i \sin(\theta_i), h^f, \hat{d}_i \cos(\theta_i))$, where h^f is the height from the camera center to the floor and θ_i is the angle of the logitude. It is noteworthy that since the expected normal vectors are parallel to the floor, thus the height of both \mathbf{p}_i and predicted

Table 1. Estimation performance based on the number of layout corners. We can see our iBARLE achieves the highest performances in almost all metrics in all corner numbers separately. It shows the effectiveness of iBARLE for improve the performances of all cases without sacrifice some specific subsets.

Corner	HorizonNet [38]				LED ² -Net [41]				LGT-Net [19]				iBARLE (Ours)			
	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1
4	86.07	84.16	0.19	0.94	86.24	84.42	0.19	0.93	87.21	85.37	0.17	0.94	88.22	86.38	0.18	0.94
5	83.76	80.66	0.28	0.92	84.49	82.27	0.25	0.92	85.76	83.44	0.22	0.93	87.83	85.74	0.20	0.93
6	83.73	81.78	0.22	0.93	83.29	81.32	0.22	0.93	83.50	81.66	0.21	0.93	85.50	83.57	0.19	0.94
7	76.02	73.54	0.30	0.90	77.45	73.86	0.29	0.89	79.68	77.22	0.28	0.91	79.62	76.92	0.25	0.92
8	80.38	78.18	0.22	0.93	79.97	77.72	0.22	0.93	80.13	77.98	0.23	0.92	80.69	78.55	0.20	0.94
9	81.92	79.59	0.25	0.92	81.65	79.17	0.25	0.91	80.39	78.17	0.26	0.92	81.14	78.75	0.23	0.93
10+	75.42	72.61	0.30	0.91	74.75	71.91	0.29	0.91	75.21	72.26	0.29	0.90	76.16	73.39	0.25	0.92
Average	81.04	78.65	0.25	0.92	81.12	78.67	0.24	0.92	81.70	79.44	0.24	0.92	82.74	80.47	0.21	0.93

Table 2. Results comparison on ZInD dataset split by type of layout. Results on each group with specific type of layout are reported separately, and the group-wise average results are also compared.

Room	HorizonNet [38]				LED ² -Net [41]				LGT-Net [19]				iBARLE (Ours)				
	Type	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1
Cuboid		86.47	84.54	0.19	0.94	86.63	84.78	0.19	0.94	87.54	85.69	0.17	0.94	88.62	86.76	0.18	0.94
Manhattan-l		83.50	81.61	0.21	0.93	83.03	81.15	0.22	0.93	83.29	81.43	0.21	0.93	85.13	83.21	0.19	0.94
Manhattan-g		78.08	76.00	0.25	0.92	77.66	75.45	0.24	0.92	78.19	75.90	0.25	0.92	78.90	76.89	0.21	0.93
non-Manhattan		79.99	76.97	0.28	0.91	80.18	77.23	0.27	0.91	80.83	78.33	0.25	0.92	82.08	79.36	0.23	0.93
Average		82.01	79.78	0.23	0.93	81.87	79.66	0.23	0.92	82.46	80.34	0.22	0.93	83.68	81.55	0.20	0.94

$\hat{\mathbf{p}}_i$ are set as h^f , which will not influence the computation of normal vectors. Then ground-truth and predicted normal vectors at the same position are computed as:

$$\begin{aligned} \mathbf{n}_i &= \mathbf{M}_r \begin{pmatrix} \mathbf{p}_{i+1} & \mathbf{p}_i \\ k\mathbf{p}_{i+1} & \mathbf{p}_i k \end{pmatrix}^\top, \\ \hat{\mathbf{n}}_i &= \mathbf{M}_r \begin{pmatrix} \hat{\mathbf{p}}_{i+1} & \hat{\mathbf{p}}_i \\ k\hat{\mathbf{p}}_{i+1} & \hat{\mathbf{p}}_i k \end{pmatrix}^\top, \end{aligned} \quad (6)$$

where \mathbf{M}_r is the rotation matrix of $\frac{\mathbf{p}_i}{2}$, and \mathbf{n}_i and $\hat{\mathbf{n}}_i$ are the ground-truth and predicted normal vectors at the same position, respectively. The learning objective is defined as maximizing inner product between the predicted and ground-truth normals as:

$$L_n = \frac{1}{N} \sum_{i=1}^N \mathbf{n}_i \cdot \hat{\mathbf{n}}_i. \quad (7)$$

Moreover, the normals change near the corners, thus the gradient of the normal angles are obtained to supervise the turning of corners [19]. However, for complex indoor space beyond simple cuboid or Manhattan world, occlusions appear and the normals near the boundaries remain consistent but the depth change sharply. In order to capture the changing of visible corners and invisible boundaries near occlusions, the gradient constraints to both the normal vectors and depth prediction are applied. Specifically, the gradient of the normals and depth are calculated as:

$$\begin{aligned} \mathbf{g}_i^n &= \arccos(\mathbf{n}_i \cdot \mathbf{n}_{i+1}), \quad g_i^d = d_{i+1} - d_i, \\ \hat{\mathbf{g}}_i^n &= \arccos(\hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_{i+1}), \quad \hat{g}_i^d = \hat{d}_{i+1} - \hat{d}_i, \end{aligned} \quad (8)$$

where $\mathbf{n}_i/\hat{\mathbf{n}}_i$ and d_i/\hat{d}_i are the ground truth and predicted normal/depth, respectively. Then, the gradient-based normal and depth prediction constraint is defined as:

$$L_g = \frac{1}{N} \sum_{i=1}^N (j\mathbf{g}_i^n \cdot \hat{\mathbf{g}}_i^n + jg_i^d \cdot \hat{g}_i^d). \quad (9)$$

To this end, the aggregated learning objective of layout estimation with the panorama images as input is obtained via integrating the aforementioned losses as $L(\mathbf{Z}) = L_d + L_h + L_n + L_g$, where \mathbf{Z} denotes the visual features extracted by $G(\cdot)$ with panorama images as input.

4.5. Overall Training Objective

Combining the horizon depth prediction, room height prediction, normals prediction, and gradient-based prediction constraint loss for all real train samples and synthesized data, produced by appearance variation domain generalization and cross-room structure mix-up, the overall learning objective of our proposed model is shown below:

$$\begin{aligned} \min_{G,F} L(\mathbf{Z}) + \alpha L(\tilde{\mathbf{Z}}) + \beta L(\bar{\mathbf{Z}}), \\ \min_{E_v, D_v} L_{var}, \end{aligned} \quad (10)$$

where α and β are hyper-parameters to balance the contribution of the real data, and the synthesized samples via AVG and CSMix, respectively. Moreover, the networks $E_v(\cdot)$ and $D_v(\cdot)$ in AVG are trained separately and fixed during the training of other networks.

Table 3. Estimation performance based on different panorama capture locations. Our iBARLE achieves the highest performances for all primary/secondary categories and the overall performance.

Camera	HorizonNet [38]				LED ² -Net [41]				LGT-Net [19]				iBARLE (Ours)			
	Pose	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE	δ_1	2DIoU	3DIoU	RMSE
Primary	85.58	83.71	0.21	0.94	85.93	84.00	0.20	0.93	86.23	84.41	0.19	0.94	87.72	85.85	0.19	0.94
Secondary	81.16	78.81	0.23	0.93	80.70	78.46	0.23	0.92	81.57	79.33	0.22	0.93	82.63	80.44	0.20	0.93
Average	83.37	81.26	0.22	0.93	83.32	81.23	0.22	0.93	83.90	81.87	0.21	0.93	85.18	83.15	0.19	0.94

Table 4. Overall experimental results comparison on ZInD dataset

Corner Number	2DIoU(%)	3DIoU(%)	RMSE #	δ_1
HorizonNet [38]	83.25	81.13	0.2219	0.9303
LED ² -Net [41]	83.18	81.08	0.2172	0.9273
LGT-Net [19]	83.81	81.77	0.2074	0.9309
iBARLE (Ours)	85.04	83.00	0.1949	0.9375

Table 5. Overall experimental results comparison on ZInD-simple

Corner Number	2DIoU(%)	3DIoU(%)	RMSE #	δ_1
HorizonNet [38]	90.44	88.59	0.123	0.957
LED ² -Net [41]	90.36	88.49	0.124	0.955
LGT-Net [ViT] [19]	88.93	86.19	0.146	0.950
LGT-Net [19]	91.77	89.95	0.111	0.960
iBARLE (Ours)	92.22	90.42	0.107	0.962

5. Experiment

5.1. Experimental settings

Datasets: Our experiments are based on two variants: (1) Zillow Indoor Dataset (ZInD) [6] is the largest indoor dataset consisting of 67,448 panorama images with room layout annotations including various and complex indoor spaces from general Manhattan, non-Manhattan, and non-flat ceilings layouts. We follow the official training/validation/test splits and adopt the “raw” layout annotations as ground truth. (2) ZInD-Simple [6] is a subset of ZInD dataset with only simple indoor cuboid layouts without any contiguous occluded corners exist. We evaluate our proposed model on ZInD-simple to compare it with the prior state-of-the-art method.

Data Splits: To evaluate iBARLE on different imbalanced subsets, we split the whole test layouts into several groups with different standards as mentioned earlier. The statistics of the splits based on corner numbers, room types, and camera poses are shown in Figure 4. We evaluate the performance of our iBARLE model for each group.

Evaluation Metrics: We use four widely used metrics for evaluation: (1) 2D IoU: Intersection over the Union of 2D room layouts. (2) 3D IoU: Intersection over Union of 3D room layouts. (3) RMSE: root mean squared error of the depth prediction with the camera height as 1.6 meters. (4) δ_1 : percentage of pixels where the ratio between the predicted depth and ground-truth depth is within a threshold of 1.25 [19, 52]. For 2DIoU, 3DIoU, and δ_1 metrics, higher

is better which is denoted by “. On the contrary, RMSE metric is a negatively-oriented score, thus lower is better denoted by #. In addition to the overall performance based on these evaluation metrics on the test data, we also report the average results across different sub-groups split by specific standards to evaluate how balanced our model on the data space across simple to complex indoor spaces.

Comparison with Baselines: We compare our layout estimation results with those of state-of-the-art baselines, namely, HorizonNet (CVPR’19) [38], LED²-Net (CVPR’21) [41], and LGT-Net (CVPR’22) [19].

5.2. Implementation Details

In our experiments, the feature extractor uses the architecture proposed in HorizonNet [38] based on ResNet-50 [14]. The architecture takes a panorama with the dimension of 512 1024 3 (height, width, channel) as input and gets 2D feature maps of 4 different scales by ResNet-50. Then, it compresses the height and up samples width N of each feature map to get 1D feature sequences with the same dimension $\mathbb{R}^{N \times \frac{D}{4}}$ and connects them, finally outputs a feature sequence $\mathbb{R}^{N \times D}$, where $D = 1024$ and $N = 256$ in our implementation. For the sequential depth prediction module, we use the SWG-Transformer proposed in [19] which is based on Transformer [40]. The whole framework is implemented with PyTorch and optimized by Adam optimizer with a learning rate set as $1e^{-4}$. For hyper-parameters, we empirically fix $\alpha = 0.1, \beta = 0.01$. The hyper-parameters of the layout estimation objectives in $L(\mathbf{Z})$ follow the same setting as [19] for a fair comparison.

5.3. Results Comparison

Layout Estimation across Imbalanced Data. The results on groups with a different number of corners and different room shapes are shown in Table 1 and Table 2, respectively. We notice the performance degradation from simple shape rooms to complex spaces, which demonstrates the motivation of exploring the imbalance issues of layout estimation across arbitrary structures. From the results, we observe that our proposed method outperforms state-of-the-art baselines in most cases with various metrics. More specifically, for the group-wise average 2D IoU and 3D IoU calculated across different numbers of corners, our model outperforms the second-best baseline by 1.04% and 1.03%, respectively.




Figure 5. Case study: Layout estimation on samples where occlusions occur and some spaces are invisible to the panorama camera. It leads to more complex, unique, and imbalanced layout shapes. “Visible” denotes the layout of visible regions. “layout-row” is the real-world room layout which includes the occlusion spaces. Our iBARLE model performs better for imbalanced/minority cases.

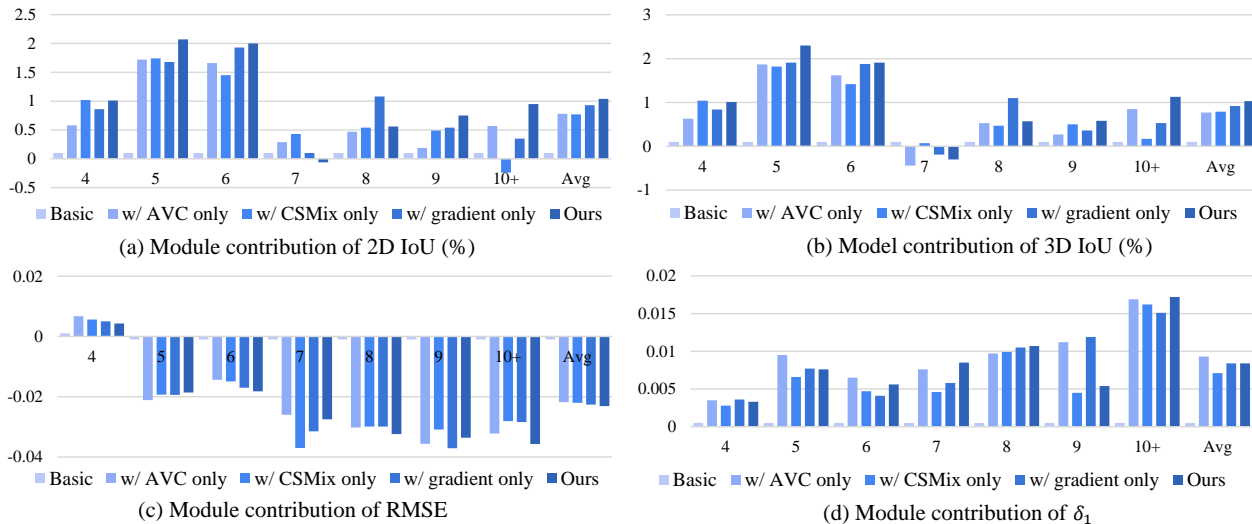


Figure 6. Ablation study of the module contributions to the iBARLE framework, where each module is separately added to the basic model. The 4 plots correspond to the 4 evaluation metrics. We can observe that each module is effective at improving the prediction performance. And in most of the cases, the complete model achieves the highest performance.

From the results in Table 2, our iBARLE model improves the 2D IoU on Manhattan-1 and non-Manhattan groups both over 1.5% compared to LGT-Net.

Moreover, we split the ZInD dataset into subsets based on the location where the camera taking the panoramas. The results are shown in Table 3. Based on the definition of the camera pose of ZInD, panoramas that are taken with “Primary” pose capture more content and are easier to capture the whole view of the room. On the contrary, images taken with the “Secondary” pose usually contain less information and are harder to estimate the layout since occlusions are more possible to occur and the camera could be too close/far from some walls. From the results, we observe our model outperforms all compared baselines on different groups and

beats the second-best baseline more than 1.0% for average 2D IoU and 3D IoU across the types of camera pose.

Overall Layout Estimation. To compare the overall layout estimation performance with prior layout estimation baselines, we report the overall results on ZInD dataset in Table 4. From the results, we observe that our proposed framework achieves a new state-of-the-art layout estimation performance for all metrics. iBARLE improves the overall 2D IoU by 1.23% over LGT-Net. Furthermore, most prior layout estimation solutions are designed for simple room shapes, e.g., cuboids. Thus, we apply our proposed model to ZInD-simple dataset, which is a subset of ZInD consisting of only simple cuboid layouts without any occluded corners exist. From the results in Table 5, we observe that our

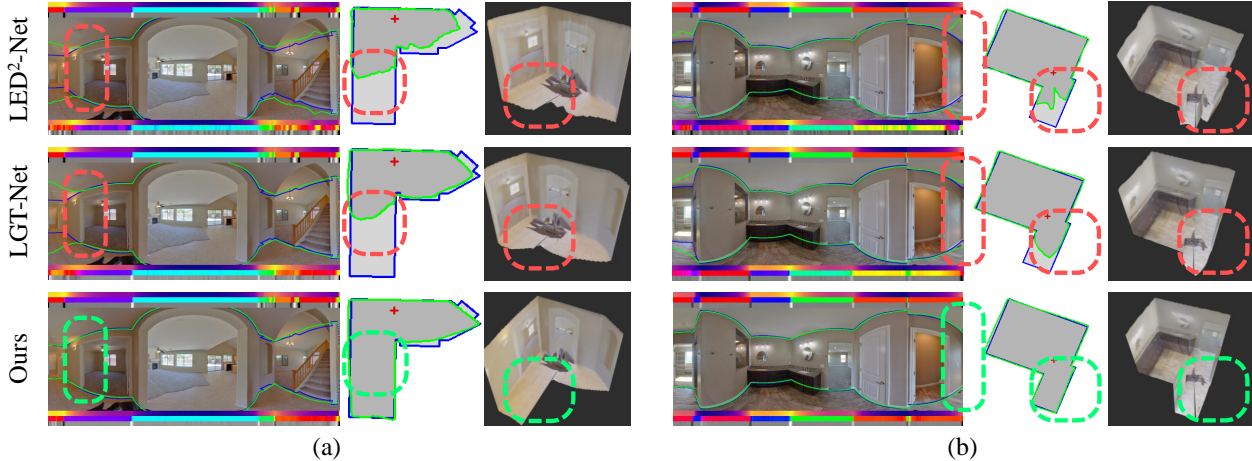


Figure 7. Case study: Layout estimation results beyond Manhattan world. Our iBARLE approach predicts non-Manhattan regions (green circle) in higher accuracy than other benchmarks. It denotes the robustness of iBARLE for handling minority cases.

model can beat the state-of-the-art baselines on simple layout subsets. Although the performance improvement on the simpler layout shape ZInD-simple subset is less significant than on the more complex and diverse ZInD dataset, the results in Table 5 demonstrate the effectiveness of the designed modules in iBARLE.

6. Discussion

Ablation Study. We further conduct an ablation study to evaluate the contribution of each module in iBARLE. Specifically, we compare the performance changes, increases of 2D IoU/3D IoU/ δ_1 and decreases of RMSE, when each of these modules is stacked to the basic model. The results changes of the four metrics are illustrated in Figure 6. From the results, we observe that all three modules are able to effectively enhance the estimation performance on almost all subgroups. It demonstrates the contribution and effectiveness of the designed three modules for layout estimation, especially for the imbalanced scenario. Moreover, our complete iBARLE model with all modules aggregated can achieve the highest performance which denotes the smoothness of the whole model structure.

Qualitative Analysis. To intuitively check the effectiveness of our iBARLE model, we visualize some results on the ZInD dataset in Figure 5 and Figure 7. The boundaries of the room layout are displayed on the panorama and the floor plan. The blue lines are the ground truth, the green lines are predictions, and the red cross is the position of the camera. In addition, the predicted horizon depth, normal, and gradient of the normal are visualized as heat maps under the panorama image, and the ground truth is shown on top of the image. Moreover, 3D visualization of the selected samples is reconstructed based on the predicted layout, and the red dash lines highlight the errors made by the compared

baselines. From the results in Figure 5, we observe that our model can manage the layout estimation of complex rooms beyond the Manhattan world assumption. Moreover, we also observe the difference between the “raw layout” and the layout “visible” to the camera. Our model can predict accurately with occlusion corners. Besides, results shown in Figure 7 are with challenging camera poses, e.g., close to the wall or corner. From the results, we can observe the proposed model is robust to predict layout with complex space panoramas taken on arbitrary positions in the room. .

Non-Manhattan Results. From Table 2, our layout estimation performance on the “Non-Manhattan” group is not as significant as other groups. Our hypothesis is that the “Non-Manhattan” group contains samples that have far more complex and diverse structures. In addition, the number and positions of the selected columns in the CSMix module are hyper-parameters influencing the complexity of the synthesized samples. These issues require more in-depth investigation on generalizability and robustness.

7. Conclusion

We propose a new technique, iBARLE, to address problems associated with data imbalance and appearance variation for single-image room layout estimation. The key components are the Appearance Variation Generalization (AVG) and Complex Structure Mix-up (CSMix) modules, which help generalize over a wide range of complex and diverse room shapes as well as scene appearance variations. We also use a gradient-based layout estimation constraint to account for occlusions in complex layouts. Experimental results show that iBARLE improves room layout estimation across different kinds of imbalanced distributions of majority and minority groups in the dataset. iBARLE also improves overall layout estimation performance.

References

- [1] Rachel KE Bellamy, Kuntal Dey, Michael Hind, Samuel C Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Jacquelyn Martino, Sameep Mehta, Aleksandra Mjilović, et al. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM Journal of Research and Development*, 63(4/5):4–1, 2019. [2](#)
- [2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Proceedings of the Advances in Neural Information Processing Systems*, 19, 2006. [3](#)
- [3] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828, 2013. [2](#)
- [4] Richard Berk, Hoda Heidari, Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. A convex framework for fair regression. *arXiv preprint arXiv:1706.02409*, 2017. [2](#)
- [5] James M Coughlan and Alan L Yuille. Manhattan world: Compass direction from a single image by bayesian inference. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2, pages 941–947, 1999. [2](#)
- [6] Steve Cruz, Will Hutchcroft, Yuguang Li, Naji Khosravan, Ivaylo Boyadzhiev, and Sing Bing Kang. Zillow Indoor Dataset: Annotated floor plans with 360 panoramas and 3D room layouts. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2133–2143, 2021. [1](#), [6](#)
- [7] Brian d’Alessandro, Cathy O’Neil, and Tom LaGatta. Conscientious classification: A data scientist’s guide to discrimination-aware classification. *Big Data*, 5(2):120–134, 2017. [2](#)
- [8] Zhengming Ding and Yun Fu. Deep domain generalization with structured low-rank constraint. *IEEE Transactions on Image Processing*, 27(1):304–313, 2017. [3](#)
- [9] Antonio D’Innocente and Barbara Caputo. Domain generalization with domain-specific aggregation modules. In *Proceedings of the German Conference on Pattern Recognition*, pages 187–198. Springer, 2018. [3](#)
- [10] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the International Conference on Machine Learning*, pages 1126–1135, 2017. [3](#)
- [11] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1180–1189, 2015. [3](#)
- [12] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(1):2096–2030, 2016. [3](#)
- [13] Thomas Grubinger, Adriana Birlutiu, Holger Schöner, Thomas Natschläger, and Tom Heskes. Domain generalization based on transfer component analysis. In *Proceedings of the Advances in Computational Intelligence*, pages 325–334. Springer, 2015. [3](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. [6](#)
- [15] Narges Honarvar Nazari and Adriana Kovashka. Domain generalization using shape representation. In *Proceedings of the European Conference on Computer Vision*, pages 666–670. Springer, 2020. [2](#)
- [16] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. [4](#)
- [17] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *Proceedings of the European Conference on Computer Vision*, pages 124–140. Springer, 2020. [3](#)
- [18] Seogkyu Jeon, Kibeom Hong, Pilhyeon Lee, Jewook Lee, and Hyeran Byun. Feature stylization and domain-aware contrastive learning for domain generalization. In *Proceedings of the ACM International Conference on Multimedia*, pages 22–31, 2021. [3](#)
- [19] Zhigang Jiang, Zhongzheng Xiang, Jinhua Xu, and Ming Zhao. LGT-Net: Indoor panoramic room layout estimation with geometry-aware transformer network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1654–1663, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [20] Xin Jin, Cuiling Lan, Wenjun Zeng, Zhibo Chen, and Li Zhang. Style normalization and restitution for generalizable person re-identification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3143–3152, 2020. [3](#)
- [21] Aditya Khosla, Tinghui Zhou, Tomasz Malisiewicz, Alexei A Efros, and Antonio Torralba. Undoing the damage of dataset bias. In *Proceedings of the European Conference on Computer Vision*, pages 158–171. Springer, 2012. [3](#)
- [22] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. SelfReg: Self-supervised contrastive regularization for domain generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9619–9628, 2021. [3](#)
- [23] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5400–5409, 2018. [3](#)
- [24] Lei Li, Ke Gao, Juan Cao, Ziyao Huang, Yepeng Weng, Xi-aoyue Mi, Zhengze Yu, Xiaoya Li, and Boyang Xia. Progressive domain expansion network for single domain generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 224–233, 2021. [2](#)
- [25] Zheren Li, Zhiming Cui, Sheng Wang, Yuji Qi, Xi Ouyang, Qitian Chen, Yuezhi Yang, Zhong Xue, Dinggang Shen, and Jie-Zhi Cheng. Domain generalization for mammography detection via multi-style and multi-view contrastive learning. In *Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 98–108. Springer, 2021. [3](#), [4](#)

- [26] Massimiliano Mancini, Samuel Rota Buló, Barbara Caputo, and Elisa Ricci. Best sources forward: Domain generalization through source-specific nets. In *Proceedings of the IEEE International Conference on Image Processing*, pages 1353–1357, 2018. [3](#)
- [27] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6):1–35, 2021. [2](#)
- [28] Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5715–5725, 2017. [3](#)
- [29] Li Niu, Wen Li, and Dong Xu. Multi-view domain generalization for visual recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4193–4201, 2015. [3](#)
- [30] Giovanni Pintore, Marco Agus, and Enrico Gobbetti. AtlantaNet: Inferring the 3D indoor layout from a single 360 image beyond the manhattan world assumption. In *Proceedings of the European Conference on Computer Vision*, pages 432–448. Springer, 2020. [1](#), [2](#)
- [31] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. [2](#)
- [32] Mohammad Mahfujur Rahman, Clinton Fookes, Mahsa Bakhtashmotlagh, and Sridha Sridharan. Multi-component image translation for deep domain generalization. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, pages 579–588, 2019. [2](#)
- [33] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *Proceedings of the International Conference on Machine Learning*, pages 1842–1850. PMLR, 2016. [3](#)
- [34] Grant Schindler and Frank Dellaert. Atlanta world: An expectation maximization framework for simultaneous low-level edge grouping and camera calibration in complex man-made environments. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, 2004. [2](#)
- [35] Yuge Shi, Jeffrey Seely, Philip Torr, Siddharth N, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. In *International Conference on Learning Representations*, 2022. [3](#)
- [36] Connor Shorten and Taghi M Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of Big Data*, 6(1):1–48, 2019. [2](#)
- [37] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Proceedings of the Advances in Neural Information Processing Systems*, 30, 2017. [3](#)
- [38] Cheng Sun, Chi-Wei Hsiao, Min Sun, and Hwann-Tzong Chen. HorizonNet: Learning room layout with 1d representation and pano stretch data augmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1047–1056, 2019. [1](#), [2](#), [5](#), [6](#)
- [39] Cheng Sun, Min Sun, and Hwann-Tzong Chen. HoHoNet: 360 indoor holistic understanding with latent horizontal features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2573–2582, 2021. [1](#), [2](#)
- [40] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. volume 30, 2017. [6](#)
- [41] Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, and Yi-Hsuan Tsai. LED2-Net: Monocular 360 layout estimation via differentiable depth rendering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12951–12960, 2021. [1](#), [2](#), [3](#), [5](#), [6](#)
- [42] Haiyan Wang, Will Hutchcroft, Yuguang Li, Zhiqiang Wan, Ivaylo Boyadzhiev, Yingli Tian, and Sing Bing Kang. PSM-Net: Position-aware stereo merging network for room layout estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8616–8625, 2022. [1](#), [2](#)
- [43] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, Tao Qin, Wang Lu, Yiqiang Chen, Wenjun Zeng, and Philip Yu. Generalizing to unseen domains: A survey on domain generalization. *IEEE Transactions on Knowledge and Data Engineering*, 2022. [2](#)
- [44] Wenhao Wang, Shengcai Liao, Fang Zhao, Cuicui Kang, and Ling Shao. DomainMix: Learning generalizable person re-identification without human annotations. 2021. [2](#)
- [45] Yufei Wang, Haoliang Li, and Alex C Kot. Heterogeneous domain generalization via domain mixup. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3622–3626, 2020. [2](#)
- [46] Guile Wu and Shaogang Gong. Collaborative optimization and aggregation for decentralized domain generalization and adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6484–6493, 2021. [3](#)
- [47] Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 14383–14392, 2021. [2](#)
- [48] Shang-Ta Yang, Fu-En Wang, Chi-Han Peng, Peter Wonka, Min Sun, and Hung-Kuo Chu. Dula-Net: A dual-projection network for estimating room layouts from a single rgb panorama. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3363–3372, 2019. [1](#), [2](#)
- [49] Yinda Zhang, Shuran Song, Ping Tan, and Jianxiong Xiao. PanoContext: A whole-room 3D context model for panoramic scene understanding. In *Proceedings of the European Conference on Computer Vision*, pages 668–686. Springer, 2014. [1](#)
- [50] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with MixStyle. In *Proceedings of the International Conference on Learning Representations*, 2020. [2](#), [4](#)

- [51] Chuhan Zou, Alex Colburn, Qi Shan, and Derek Hoiem. LayoutNet: Reconstructing the 3D room layout from a single RGB image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2051–2059, 2018. [2](#)
- [52] Chuhan Zou, Jheng-Wei Su, Chi-Han Peng, Alex Colburn,

Qi Shan, Peter Wonka, Hung-Kuo Chu, and Derek Hoiem. Manhattan room layout reconstruction from a single 360 image: A comparative study of state-of-the-art methods. *International Journal of Computer Vision*, 129(5):1410–1431, 2021. [2](#), [6](#)