

## U2RLE: Uncertainty-Guided 2-Stage Room Layout Estimation

Pooya Fayyazsanavi<sup>\*†1</sup>, Zhiqiang Wan<sup>\*2</sup>, Will Hutchcroft<sup>2</sup>, Ivaylo Boyadzhiev<sup>2</sup>, Yuguang Li<sup>2</sup>, Jana Kosecka<sup>1</sup>, Sing Bing Kang<sup>2</sup>

<sup>1</sup>George Mason University <sup>2</sup>Zillow Group

### Abstract

While the existing deep learning-based room layout estimation techniques demonstrate good overall accuracy [17], they are less effective for distant floor-wall boundary. To tackle this problem, we propose a novel uncertainty-guided approach for layout boundary estimation introducing new two-stage CNN architecture termed U2RLE. The initial stage predicts both floor-wall boundary and its uncertainty and is followed by the refinement of boundaries with high positional uncertainty using a different, distance-aware loss. Finally, outputs from the two stages are merged to produce the room layout. Experiments using ZInD [4] and Structure3D [25] datasets show that U2RLE improves over current state-of-the-art, being able to handle both near and far walls better. In particular, U2RLE outperforms current state-of-the-art techniques for the most distant walls.

### 1. Introduction

A lot of progress has been made in image-based room layout estimation (RLE). While deep learning techniques applied to RLE have been very effective, they are evaluated using datasets of rooms that are mostly small and cuboid-shaped [1, 23]. However, ZInD [4], which is a large-scale dataset of real homes, shows that real rooms are large and have more complex layouts.

The challenge of distant wall depth estimation can be seen clearly from the relationship between pixel and depth errors in Fig. 1. We simulate a perturbed layout by shifting the vertical position of each room corner of the GT layout by 3 pixels. With an image height of 512 pixels, 3 pixels corresponds to 0.59% relative error in image-space, for both near and distant walls, alike; however, this results in significantly larger depth error for distant walls. For the room with

<sup>\*</sup>Equal contribution.

<sup>†</sup>This work was done when Pooya Fayyazsanavi was an intern at Zillow.

<sup>1</sup>{pfayyazs,kosecka}@gmu.edu

<sup>2</sup>{zhiqiangw,willhu,ivaylob,yuguangl,singbingk}@zillowgroup.com

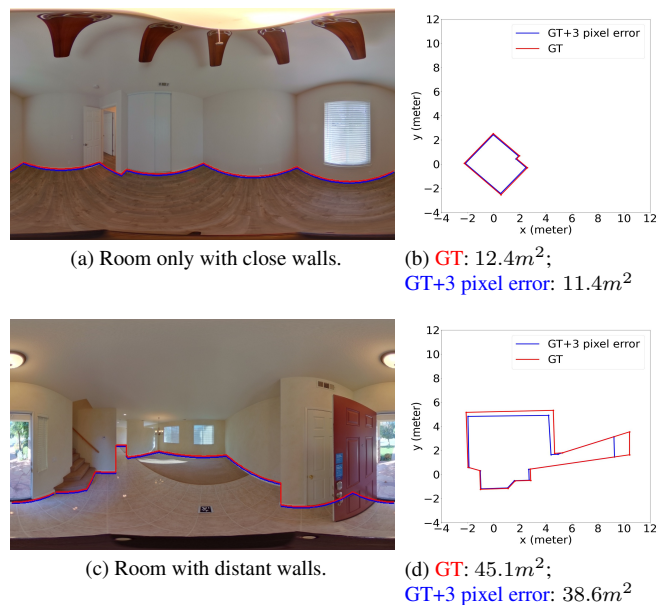


Figure 1. The same pixel error from a floor point maps to a larger depth error when the point is further away from the camera. GT layout is in Red while the simulated perturbed layout is in Blue. (b) and (d) share aligned coordinate system.

only close walls (Fig. 1b), the area of the GT layout and the perturbed layout are  $12.4m^2$  and  $11.4m^2$ , respectively, resulting in just  $1m^2$  error. For a room with distant walls, the area of the GT layout and the perturbed layout are  $45.1m^2$  and  $38.6m^2$ , respectively, an error of  $6.5m^2$ . As we can see, distant walls have a big impact on estimating the area of a room. For real-world applications such as real estate, where the floorplan size is a major factor in pricing, this magnitude of the absolute error is undesirable, pointing to the need for new methods to address this challenge.

While recent large-scale datasets [4, 25] provide large annotated spaces with more distant walls, dataset imbalance challenges accurate estimation. For example, in ZInD [4],

nearly 90% of the walls are within 4 meters of the camera while the remaining walls can extend out to 10+ meters. Many current models are trained to minimize the loss averaged across all image columns; imbalance may lead to a model that does not prioritize the feature resolution and granularity necessary to accurately estimate distant walls.

To address these challenges, we propose a novel two-stage architecture. The first (initial) stage is designed to predict the wall position, alongside a measure of uncertainty. The second (refine) stage focuses on distant regions using distance-aware loss. Finally, guided by the predicted uncertainty score, the outputs of these two stages are merged to generate the room layout. This two-stage architecture allows us to design specific loss functions and data augmentation for close vs. distant walls; it handles the imbalance in the wall depth distribution.

Our contributions are:

- Uncertainty-guided two-stage (initial and refine) room layout estimation model, achieving SOTA performance on multiple datasets.
- Novel loss function to predict floor-wall boundary uncertainty in the initial stage.
- New Channel-Preserving Height Compression (CPHC) module that compresses features along its height.
- New distance-aware loss to re-weight the influence of distant walls in the refinement stage.

## 2. Related Work

**Single-view Room Layout Estimation.** Single-view room layout estimation has been an active area of research for the past decade [11]. A good amount of work has been done to generate partial room layouts from a single perspective image. Learning-based approaches [9, 14, 21] have been proven more effective compared to geometry-based ones [5, 6]. PanoContext [23] is the first method to demonstrate the effectiveness of room layout estimation using a single 360° equirectangular panorama image. It generates the room layout by estimating per-pixel normals of overlapping perspective images converted from a single panorama. LayoutNet [26] demonstrated the benefits of using encoder-decoder CNN-based architecture to directly operate on the panoramic image; layouts are predicted by estimating the probability maps of room boundary map and corners, followed by a 3D layout optimization.

HorizonNet [17] simplified layout estimation tasks by generating a 1D representation of the floor-wall, ceiling-wall, and wall-wall boundary positions per image column; it uses a bidirectional LSTM (Bi-LSTM) module to regress across the panorama. Since then, many approaches are

based on the HorizonNet architecture. HoHoNet [18] re-designed the feature extractor with the Efficient Height Compression (EHC) module followed by multi-head self-attention (MSA) to generate Latent Horizontal Feature (LHFeat). They demonstrated good performance using LHFeat to model per-column modality for room layout reconstruction as well as dense predictions such as depth estimation and semantic segmentation. LED2Net [19] adopted HorizonNet’s backbone and applied differentiable depth rendering to incorporate the room polygon geometry information into their end-to-end model. LGT-Net [7] used a novel transformer architecture as their sequence processor to improve the spatial identification ability for the panorama. LED2Net and LGT-Net both achieve state-of-the-art performance overall; the recent focus of layout estimation study has placed more emphasis on solving complex, non-Manhattan room shapes, which are found more frequently in real-life, large-scale datasets such as ZinD [4] and CG datasets such as Structure3D [25].

While using equirectangular panorama images as inputs, DuLaNet [22], AtlantaNet [13], and PSMNet [20] also leverage the Equirectangular-to-Perspective (E2P) transform to generate a binary segmentation map as the floor plan. AtlantaNet demonstrated an advantage of modeling non-Manhattan and complex room layouts. DMHT [24] alternatively applied a learnable Hough Transformation Block on cubemap tiles from panoramic images and claimed SOTA performance.

**Multi-view Room Layout Estimation.** Additional viewpoints can be used to improve layout estimation for complex, large rooms [16, 20]. Furthermore, recent work [15] has used multi-view consistency, in a self-supervised approach, to improve single-room layout estimation.

**Uncertainty Estimation.** Several approaches for depth estimation in the past used iterative techniques to improve the depth estimates either using additional cues, such as plane normals [2]. General framework for uncertainty estimation in Deep Learning has also been adopted in the past for geometric regression problems [10]. These class of methods uses Monte-Carlo (MC) sampling to estimate uncertainty, where the samples are generally computed using an ensemble of neural networks. The application of uncertainty to the problem of estimating room layouts is however new. Here we use simpler yet effective approach where both the mean and the variance are directly predicted from data.

## 3. System Overview

We present U2RLE, a novel two-stage approach to estimate the layout of a room using a single 360° equirectangular panorama. Figure 2 shows the overall architecture of

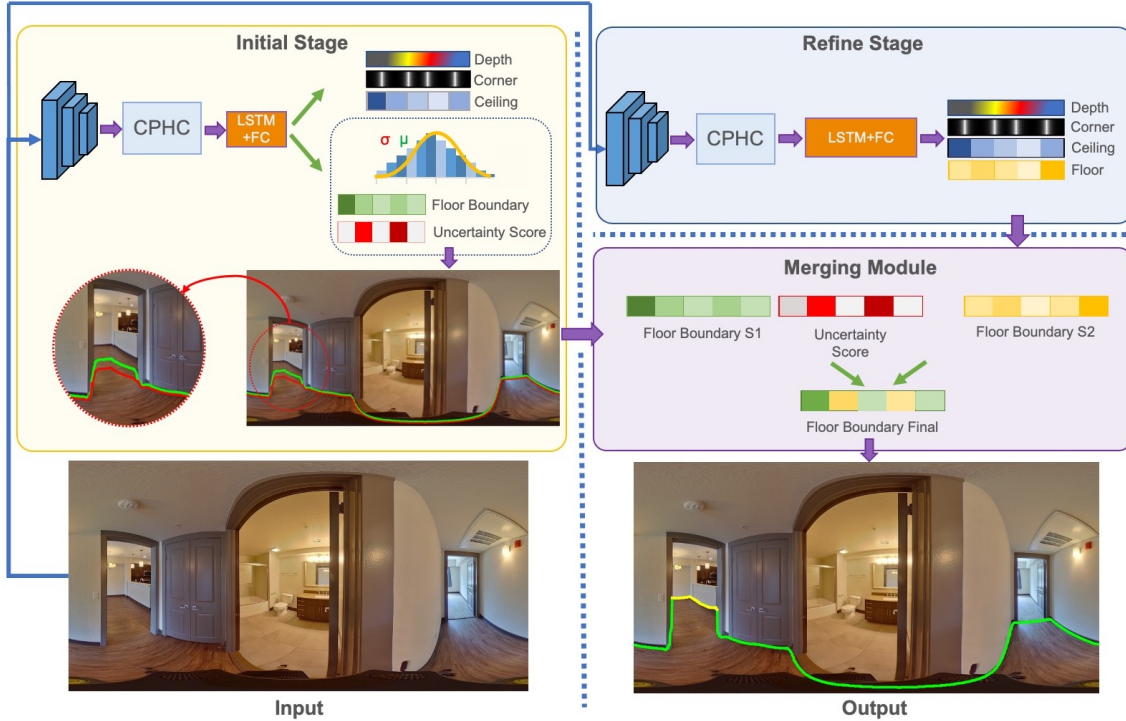


Figure 2. Overview of U2RLE. It has three main components: initial stage, refinement stage, and merging module. In the initial stage, the model predicts the boundary and estimates its positional uncertainty. In the initial stage, the gap between the green boundary and red boundary represents the uncertainty score. The refinement stage focuses on distant regions using distance-aware loss. Finally, the merging module combines the predictions from both stages using the uncertainty score. In the merging module, the uncertain parts indicating by red will be replaced by the floor-wall boundary prediction from the refinement stage.

our method. We describe the initial stage in section 4. The goal of the initial stage is to provide initial boundary and uncertainty predictions. The goal of the refinement stage, which we describe in section 5, is to improve the far away region using a novel distance-aware loss function. Finally, in section 6, we discuss how to utilize the predicted uncertainty score to combine predictions from both stages to get the best of both worlds, i.e. close and distant regions.

#### 4. Initial Boundary and Uncertainty Prediction

The initial stage follows the approach of HorizonNet [17]. The input image of size  $512 \times 1024 \times 3$  is passed through ResNet-50 to extract features. Then, we propose a new Channel-Preserving Height Compression Module (CPHC) to compress the information along height dimension, so as to obtain per column features that can be used for boundary prediction. The compressed features in the column order are then passed through bi-directional LSTM and fully-connected layers to predict the boundary and the uncertainty score for each column of the panorama, along with corners and depth.

#### 4.1. Channel-preserving Height Compression

HorizonNet [17] proposed the height compression module (HCM) to effectively squeeze features with 2D spatial support from the backbone along the height dimension to produce a 1D horizontal feature. Many recent room layout approaches [3, 7, 19] adopt HCM yielding good results. The features from block 1 ~ 4 of ResNet-50 are compressed to dimensions (channel  $\times$  height  $\times$  width)  $32 \times 8 \times 256, 64 \times 4 \times 256, 128 \times 2 \times 256, 256 \times 1 \times 256$ . As we can see, the height dimension of block 1 ~ 3 has not been compressed to 1. In HCM, the first two dimensions (channel and height) will be reshaped into one dimension. We argue this is not the most effective way to compress the height information. Since we predict one value for each image column, it is better not to mix the height and channel dimensions. To address this issue, we propose a new Channel-Preserving Height Compression (CPHC) Module that compresses the height dimension into 1 and avoids mixing the information from feature channels and height dimension. The architecture of the proposed CPHC is shown in supplementary material. The ablation study in Tab. 4 demonstrates CPHC module’s ability to enhance the performance of both

initial and refinement stage.

## 4.2. Uncertainty Prediction

Unlike the existing models, the initial stage not only predicts the floor-wall boundary but also the uncertainty of the prediction. Specifically, the model predicts two quantities per each column namely floor-wall boundary mean  $\mu$ , and floor-wall boundary standard deviation  $\sigma$ .

The model is optimized by minimizing the negative log likelihood loss function for each image column  $i$  with  $\mu_i$  and  $\sigma_i$ :

$$\mathcal{L}_{floor}^{NL} = - \sum_{i=1}^W \log \frac{1}{\sigma_i \sqrt{2\pi}} \exp \left( - \frac{(y_i - \mu_i)^2}{2(\sigma_i)^2} \right) \quad (1)$$

where  $y_i$  is the ground truth floor-wall boundary, and  $W$  is the image width. During the inference, the column-wise predicted mean  $\mu_i$  vector will be used as the predicted floor-wall boundary. Since the predicted uncertainty vector  $\sigma_i$  is in the image-space, we need to project it to the depth-space. We begin by calculating the *upper bound*  $\hat{y}$  of the floor-wall boundary by adding the predicted mean vector  $\mu$  to predicted standard deviation vector  $\sigma$ :

$$\hat{y} = \mu + \sigma \quad (2)$$

The predicted floor-wall boundary ( $\mu$ ) and the upper bound  $\hat{y}$  are projected into the depth-space omitting the index  $i$  for clarity. We define the uncertainty score as the distance between the projected floor-wall boundary and the upper bound.

$$U_{floor} = |Proj(\hat{y}) - Proj(\mu)| \quad (3)$$

where  $Proj(\cdot)$  is the projection function to depth. Details about this projection function can be found in [7].

## 4.3. Loss Function

The total loss is calculated as

$$\mathcal{L} = \mathcal{L}_{floor}^{NL} + \mathcal{L}_{ceil} + \mathcal{L}_{depth} + \lambda \mathcal{L}_{corner} \quad (4)$$

where the negative likelihood loss  $\mathcal{L}_{floor}^{NL}$  of floor-wall boundary is described in Eq. 1;  $\mathcal{L}_{ceil}$  is  $\ell_1$  loss of ceiling-wall boundary;  $\mathcal{L}_{depth}$  represents the  $\ell_1$  loss of wall depth;  $\mathcal{L}_{corner}$  is the corner loss.

Unlike HorizonNet where Binary Cross-Entropy Loss is used as corner loss, we propose to replace it with a  $\ell_1$  loss. This is because the ground truth is a continuous value instead of a binary value.

## 5. Fine-grained Distance-focused Prediction

Similar to the initial stage, ResNet-50 is used to extract features and the proposed CPHC compresses the extracted features. Then, bi-directional LSTM and fully-connected layers are used to predict the floor-wall boundary, the depth, the ceiling-wall boundary, and the corners.

## 5.1. Data Augmentation

In the refinement stage, we intend to focus on distant boundaries. Since these regions are located far from the camera, they are usually low resolution and blurry. The top-down view further magnifies prediction errors for these regions in comparison to close walls, as shown in Figure 1. Furthermore, in all current datasets the distant boundaries are among the least representative samples.

To enable the model to train on more of these samples, we first perform pano-stretch data augmentation [17] and stretch each sample’s layout to a greater distance during training. As opposed to the initial stage, which pushes the boundary points to both the close and far regions, we only use the data augmentation to stretch the layout to represent the far regions. In addition, different parameters were used to magnify the stretch. As with [7, 17, 19], we also only augment the data on the x and z axes only:

$$\begin{aligned} x' &= k_x \cdot x = k_x \cdot d \cdot \cos(v) \cdot \cos(u) \\ z' &= k_z \cdot z = k_z \cdot d \cdot \cos(v) \cdot \sin(u) \end{aligned} \quad (5)$$

where  $x, z$  is the 3D coordinate projected on top down-view from panorama image coordinate  $u, v$ .  $d$  is the depth.  $k_x$  and  $k_z$  is a random value picked from a uniform distribution of range [1, 2.5].

## 5.2. Distance-aware Loss Function

To force the model to focus more on the distant walls, we propose a distance-aware loss which contains two separate loss functions. The first loss function emphasizes depth prediction on projected top-down view as follows:

$$\mathcal{L}_{depth} = \sum_{i=1}^W \left| \hat{d}_i - d_i \right| \quad (6)$$

where  $\hat{d}$  is the predicted depth and  $d$  is the ground truth depth.

To focus more on the distant regions we also re-weight the floor boundary loss based on the distance as follow:

$$\mathcal{L}_{floor} = \sum_{i=1}^W |\hat{y}_i - y_i| * d_i \quad (7)$$

where  $\hat{y}$  represents the predicted floor-wall boundary while  $y$  is the ground truth floor-wall boundary.

For the refinement stage, the total loss function is calculated as follows:

$$\mathcal{L} = \mathcal{L}_{floor} + \nu \mathcal{L}_{depth} + \lambda \mathcal{L}_{corner} + \mathcal{L}_{ceil} \quad (8)$$

where  $\lambda, \nu \in \mathbb{R}$  are hyper-parameters designed to balance the effect of each components. Similar to the initial stage, ceiling loss  $\mathcal{L}_{ceil}$  and corner loss  $\mathcal{L}_{corner}$  use  $\ell_1$  loss.

GT Depth (m) \ Model	1	2	3	4	5	6	7	8	9	10	2DIoU
HorizonNet [17]	0.035	0.048	0.073	0.136	0.256	0.367	0.540	0.757	0.964	1.267	89.80%
LED2-Net [19]	0.036	0.049	0.070	0.125	0.222	0.343	0.526	0.800	0.903	1.156	89.95%
LGT-Net [7]	0.035	0.051	<b>0.067</b>	<b>0.123</b>	0.233	0.326	0.501	0.693	0.858	1.145	90.82%
Ours(U2RLE)	<b>0.034</b>	<b>0.047</b>	0.072	0.127	<b>0.229</b>	<b>0.313</b>	<b>0.437</b>	<b>0.614</b>	<b>0.776</b>	<b>0.942</b>	<b>91.39%</b>

Table 1. Mean depth error (m) ↓ on ZInD dataset.

GT Depth (m) \ Model	1	2	3	4	5	6	7	8	9	10	2DIoU
HorizonNet [17]	0.028	0.036	0.073	0.148	0.245	0.344	0.688	0.891	1.423	1.443	92.63%
LED2-Net [19]	0.038	0.045	0.086	0.158	0.233	0.298	0.485	<b>0.527</b>	1.152	1.053	92.15%
LGT-Net [7]	0.034	0.039	0.074	0.142	0.232	<b>0.257</b>	0.516	<b>0.558</b>	1.073	1.155	92.81%
Ours(U2RLE)	<b>0.027</b>	<b>0.031</b>	<b>0.059</b>	<b>0.128</b>	<b>0.217</b>	0.268	<b>0.505</b>	0.575	<b>0.992</b>	<b>0.877</b>	<b>93.73%</b>

Table 2. Mean depth error (m) ↓ on Structure3D dataset.

## 6. Uncertainty-guided Merging

In order to determine the final floor-wall boundary, we considered the model predictions as prior information for the merging module. The predicted uncertainty score  $\sigma$  and estimated distance (from initial floor boundary) are used to determine the output of the merging module. As described in section 4.2, the initial stage uncertainty scores can be projected into the top-down view. Projected uncertainty scores indicate the level of uncertainty in each column. Therefore, the refinement stage results can be merged based on the score. According to our empirical analysis, we used 0.2 as the threshold for uncertainty  $U_{floor}$  across all the experiments. Also, to further refine the predictions we only combine the results if the predicted wall is distant. Consequently, fewer outliers (close walls) will be merged from the refinement stage into the final prediction. We therefore use the estimated distance as a second parameter, which will be merged if the distance is greater than 5 meters. The module finally outputs the final prediction with size of  $W$ .

## 7. Experiments

We report results of our approach on Zillow Indoor Dataset (ZInD) [4] and Structure3D [25] dataset. We also provide information regarding the training schema, used datasets, and our baselines.

### 7.1. Training

We implement U2RLE-Net using PyTorch [12]. The Adam optimizer [8] is employed to train our network with  $\beta_1 = 0.9, \beta_2 = 0.999$ . The network is trained on four NVIDIA Tesla V100 GPUs for 60 epochs on ZInD dataset and 50 epochs on Structure3D dataset, with a batch size of 40. In addition, we set hyper-parameters in Eq. (4 and 8) as  $\lambda = 10, \nu = 5$ .

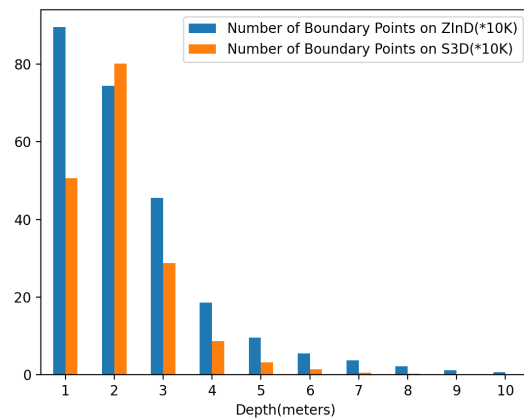


Figure 3. Depth distribution of the GT boundary points on ZInD [4] and Structure3D [25].

Apart from the pano-stretching data augmentation described in section 5 for the refinement stage, we use the same data augmentation methods proposed in Horizon-Net [17], including left-right flipping, horizontal rotation, luminance changes, and pano-stretching during training.

### 7.2. Datasets

Zillow indoor dataset (ZInD) [4] contains 71474 panoramas of 1524 unfurnished homes. It's currently the largest real world dataset with room layout annotations. It contains a close-to real world distribution of room complexities, which include a good amount of non-Manhattan walls, a wide ranges of room sizes and wall-wall corners per room. We split the dataset for training, validation, test at 20077, 2458, 2458 unique panoramas. The split is based on Horizon-Net [17] pre-processing.

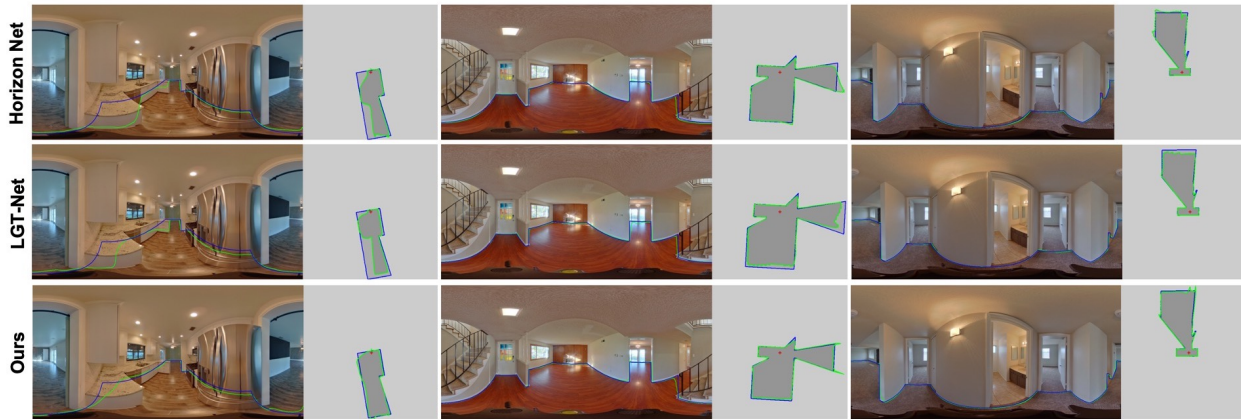


Figure 4. Qualitative comparison on ZInD [4] dataset. GT layout is in blue while predicted layout is in green.

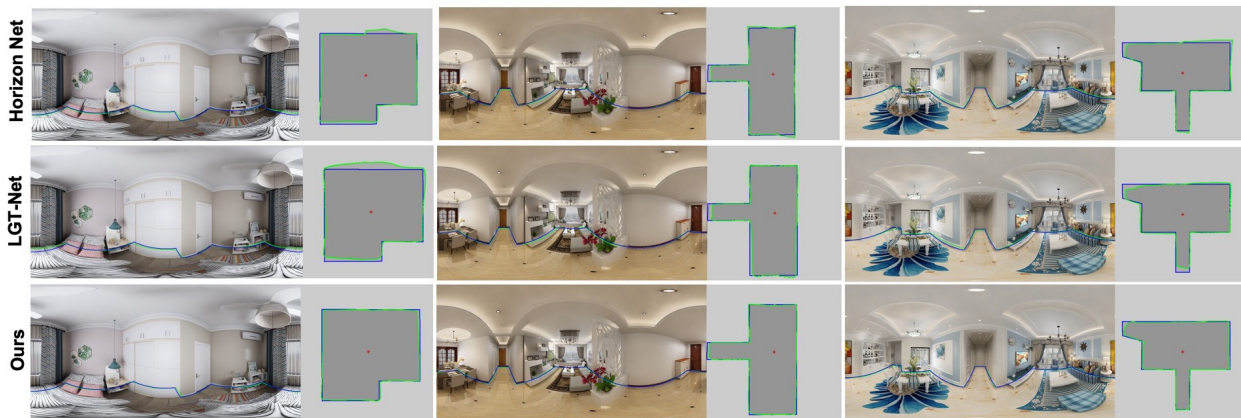


Figure 5. Qualitative comparison on Structure3D [25] dataset. GT layout is in blue while predicted layout is in green.

Structure3D [25] contains more than 20K synthetic images from different rooms with 3D room layout annotations. Structure3D contains a good range of room complexity and it is furnished. The full dataset was used. The first 3000 scenes will be used for training, 250 scenes for validation and the remaining 250 scenes for testing.

To better understand the imbalanced issue of the dataset, we plot the depth distribution of the boundary points on ZInD and Structure3D. The boundary point is obtained by projecting the floor-wall boundary point from image-space onto depth-space. Details about this projection can be found at section 4.2. As we can see from Fig. 3, ZInD and Structure3D are highly imbalanced. For example, in ZInD, around 90% of the boundary points are within 4 meters of the camera while the depth of the remaining boundary points can extend to 10+ meters.

### 7.3. Baselines

Currently, none of the models have been trained on the "visible geometry" of ZInD. Therefore, we retrained all the models based on their repository on the ZinD dataset. Additionally, if needed, baselines are also trained on Structure3D. We compare our U2RLE with baselines built upon recent state-of-the-art (SOTA) layout estimation methods: HorizonNet [17], LED2-Net [19] and, LGT-Net [7]. Also, due to the complexity of the structures in "visible geometry", post processing techniques introduced in [17, 22] do not perform effectively and can significantly drop the performance. Therefore, for evaluation purposes, we evaluate all the models without the post-processing techniques.

### 7.4. Evaluation Metrics

In order to show the effectiveness of our model, we reported our error using two metrics, mean depth error and

GT Depth (m) \ Model	1	2	3	4	5	6	7	8	9	10
Initial Stage	<b>0.034</b>	<b>0.045</b>	<b>0.067</b>	<b>0.118</b>	0.223	0.330	0.514	0.790	0.999	1.430
Refinement Stage	0.073	0.063	0.080	0.124	<b>0.215</b>	<b>0.278</b>	<b>0.397</b>	<b>0.540</b>	<b>0.705</b>	<b>0.854</b>
Two-stage (U2RLE)	0.034	0.047	0.072	0.127	0.229	0.313	0.437	0.614	0.776	0.942

Table 3. Ablation study: two-stage vs single-stage.

GT Depth (m) \ Model	1	2	3	4	5	6	7	8	9	10
Initial Stage	<b>0.034</b>	<b>0.045</b>	<b>0.067</b>	<b>0.118</b>	<b>0.223</b>	<b>0.330</b>	<b>0.514</b>	<b>0.790</b>	<b>0.999</b>	<b>1.430</b>
Initial Stage Without CPHC module	0.042	0.055	0.080	0.153	0.291	0.474	0.724	1.121	1.151	2.086
Refinement Stage	0.073	<b>0.063</b>	0.080	<b>0.124</b>	<b>0.215</b>	<b>0.278</b>	<b>0.397</b>	<b>0.540</b>	<b>0.705</b>	<b>0.854</b>
Refinement Stage Without CPHC module	<b>0.067</b>	0.066	<b>0.078</b>	0.128	0.227	0.315	0.458	0.622	0.785	1.056

Table 4. Ablation study: proposed CPHC module.

2D IoU. 2D IoU shows the overall performance of the models, but it does not explicitly show the effectiveness of predictions based on distance. As shown in Figure 1, the error on the distant regions will be amplified and can degrade 2D IoU significantly. To better study the prediction errors, we evaluate all of the models based on the mean depth error. To calculate the error, first the predicted and ground truth boundary are projected to top-down view. The distance from the projected ground truth boundary to the origin is used to calculate the error. The difference between the predicted and ground truth boundary is defined as the depth error. To show the overall model’s performance, we also report 2D IoU.

### 7.5. Quantitative Evaluation

Our evaluation on ZInD [4], is shown in Table 1. All the errors are reported in meters. Our approach has lower mean depth error on distant regions in comparison to all other approaches. Also, the model can outperform the available models in the overall 2D IoU. To show the effectiveness of our model in furnished settings, we compare our approach to other baselines. Table 2 shows the introduced metrics on Structure3D [25] dataset. Our approach can outperform significantly on almost all close and distant regions.

### 7.6. Qualitative Evaluation

Figure 4 shows some examples on ZInD dataset [4] and the predicted boundary both on panorama and top-down view. The green color represents the predicted floor boundary and blue indicates the Ground Truth. The first image is an example where our model is capable of capturing both distant and close regions appropriately. The second image illustrates the effectiveness of our model in large open spaces. Finally, the last image represents a complex structure that our model is able to predict more accurately.

Figure 5 shows the results on Structure3D [25] dataset.

Compared to ZinD [4], Structure3D is furnished. First image shows the results of our model compared to other baselines when occlusions occur in close regions, where our model is able to handle the occlusion more effectively. Additionally, we provided examples of large open areas in the second and third images.

## 8. Ablation Study

We conduct experiments to determine how individual components affect the U2RLE architecture. All the ablation studies are conducted on ZInD dataset. In particular, we examine the following variants:

- Compare the proposed two-stage model with single-stage model.
- Replace the proposed Channel-Preserving Height Compression (CPHC) Module (Section 4.1) with the original height compression module proposed by [17].
- Replace the uncertainty prediction in the initial stage with  $\ell_1$  loss function to predict the boundaries directly.
- Finally, we do experiments on the effect of our distance-aware loss (Section 5.2) in the refinement stage.

### 8.1. Two-stage VS Single-stage

The motivation of designing our two-stage model is that current single stage models cannot achieve optimal performance for both close walls and distant walls. As we can see from Tab. 3, the initial stage model works well for close walls where depth ranges from 1m to 4m while while the refinement stage performs better on distant walls where depth ranges from 5m to 10m. With our proposed two-stage architecture (U2RLE), we can merge the results from the initial

GT Depth (m) \ Model	1	2	3	4	5	6	7	8	9	10
Initial Stage	<b>0.034</b>	<b>0.045</b>	<b>0.067</b>	<b>0.118</b>	<b>0.223</b>	<b>0.330</b>	0.514	0.790	0.999	1.430
Without Uncertainty	0.035	0.046	0.069	0.123	0.233	0.332	<b>0.495</b>	<b>0.703</b>	<b>0.926</b>	<b>1.242</b>

Table 5. Ablation study: uncertainty prediction.

GT Depth (m) \ Model	1	2	3	4	5	6	7	8	9	10
Refinement Stage	0.073	0.063	<b>0.080</b>	<b>0.124</b>	<b>0.215</b>	<b>0.278</b>	<b>0.397</b>	<b>0.540</b>	<b>0.705</b>	<b>0.854</b>
Without Distance-aware Loss	<b>0.073</b>	<b>0.063</b>	0.081	0.135	0.240	0.354	0.546	0.746	0.929	1.444

Table 6. Ablation study: distance-aware loss.

stage and the refinement stage to achieve good performance for both close walls and distant walls.

One may argue that we can merge based on a depth threshold like 4.5m. In reality, we do not know the GT depth. The predicted depth is pretty noisy. Directly merging based on the predicted depth will result in sub-optimal result. In our merging module, we use the predicted uncertainty score to guide the merging and utilize the predicted depth to reduce outlier, which gives better result.

## 8.2. Proposed CPHC Module

In our U2RLE, a CPHC module is proposed for the initial stage and the refinement stage to compress information along height dimension. As we can see from Tab. 4, the mean depth error of the initial stage becomes worse for all depths if we replace the proposed CPHC module with the original height compression module. Similarly, the proposed CPHC module helps to achieve better depth estimation for the refinement stage, especially for the distant walls. These results verify the effectiveness of the proposed height compression module.

## 8.3. Uncertainty Prediction

The initial stage predicts both the floor-wall boundary and uncertainty score. The uncertainty prediction is not only important for the merging step but also can improve the performance. As we can see from Tab. 5, the uncertainty prediction helps to improve the performance for close walls. One may also notice that our initial stage with uncertainty prediction has worse results for distant walls. Since the initial stage is designed to predict close walls, this will not affect the final result. This also supports our motivation of the two-stage model where we can use different loss functions for the initial stage and the refinement stage to handle close walls and distant walls.

The uncertainty prediction is optimized via negative log likelihood loss, which can be minimized in two ways for far away regions, by predicting an accurate boundary ( $\mu$ ) or

by increasing the uncertainty ( $\sigma$ ) for those far away regions. Our model has decided to minimize the loss by increasing the uncertainty there, which means that the optimizer do not need to find an accurate  $\mu$  for far away regions. In this way, the initial stage can allocate more learning capacity on close walls.

## 8.4. Distance-aware Loss

In the refinement stage, a distance-aware loss is proposed to force the model focus more on the distant walls. As we can see from Tab. 6, the proposed distance-aware loss significantly reduces the mean depth error for distant walls.

## 9. Conclusions and Future Work

We presented effective improvement of the state-of-the-art in 1D layout estimation from a single panoramic view. With similar motivation to traditional depth estimation methods from perspective views, which suffers from lower quality estimates for far away regions, our approach focused on improvements for distant boundaries. We proposed a novel CPHC model to improve the overall accuracy with floor-wall boundary confidence, introduced a refined stage to produce higher accuracy for distance regions and apply the uncertainty guided merging module for final predictions. We presented comprehensive evaluation along the ablation study showing the effect of different improved components.

Current generation of high-resolution panoramas can be easily 8K at full resolution. It is computationally expensive to use the full resolution for a single stage network and the reduction of resolution affects more the quality of the estimates in far away regions. We believe that the two-stage approach can provide benefit in this setting; lower resolution data can be used for initial stage, and the higher resolution is used only for areas with high uncertainty. Testing the effectiveness of the proposed model in the multi-resolution setting can be explored in the future work.



## References

- [1] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. **1**
- [2] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Irondepth: Iterative refinement of single-view depth using surface normal and its uncertainty. In *DOI: 10.48550/arXiv.2210.03676*. **2**
- [3] Kefan Chen, Noah Snaveley, and Ameesh Makadia. Wide-baseline relative camera pose estimation with directional learning. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3257–3267, 2021. **3**
- [4] 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/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2133–2143, June 2021. **1, 2, 5, 6, 7**
- [5] Alex Flint, Christopher Mei, David Murray, and Ian Reid. A dynamic programming approach to reconstructing building interiors. In *European conference on computer vision*, pages 394–407. Springer, 2010. **2**
- [6] Varsha Hedau, Derek Hoiem, and David Forsyth. Recovering the spatial layout of cluttered rooms. In *2009 IEEE 12th international conference on computer vision*, pages 1849–1856. IEEE, 2009. **2**
- [7] 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 (CVPR)*, 2022. **2, 3, 4, 5, 6**
- [8] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015. **5**
- [9] Chen-Yu Lee, Vijay Badrinarayanan, Tomasz Malisiewicz, and Andrew Rabinovich. Roomnet: End-to-end room layout estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 4865–4874, 2017. **2**
- [10] Antonio Loquercio, Mattia Segù, and Davide Scaramuzza. A general framework for uncertainty estimation in deep learning. In *IEEE Robotics and Automation Letters*, Vol. 5, No. 2. **2**
- [11] Bincy P Mathew. Review on room layout estimation from a single image. *International Journal of Engineering Research and*, 9, 2020. **2**
- [12] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. *ArXiv*, abs/1912.01703, 2019. **5**
- [13] Giovanni Pintore, Marco Agus, and E. Gobbetti. Atlantantet: Inferring the 3D indoor layout from a single 360-deg image beyond the Manhattan world assumption. In *ECCV*, 2020. **2**
- [14] Yuzhuo Ren, Shangwen Li, Chen Chen, and C-C Jay Kuo. A coarse-to-fine indoor layout estimation (cfile) method. In *Asian conference on computer vision*, pages 36–51. Springer, 2016. **2**
- [15] Bolivar Solarte, Chin-Hsuan Wu, Yueh-Cheng Liu, Yi-Hsuan Tsai, and Min Sun. 360-mlc: Multi-view layout consistency for self-training and hyper-parameter tuning. In *Advances in Neural Information Processing Systems*, 2022. **2**
- [16] Jheng-Wei Su, Chi-Han Peng, Peter Wonka, and Hung-Kuo Chu. Gpr-net: Multi-view layout estimation via a geometry-aware panorama registration network. *CoRR*, abs/2210.11419, 2022. **2**
- [17] Cheng Sun, Chi-Wei Hsiao, Min Sun, and Hwann-Tzong Chen. Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. **1, 2, 3, 4, 5, 6, 7**
- [18] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Hohonet: 360 indoor holistic understanding with latent horizontal features. In *CVPR*, 2021. **2**
- [19] Fu-En Wang, Yu-Hsuan Yeh, Min Sun, Wei-Chen Chiu, and Yi-Hsuan Tsai. Led2-net: Monocular 360deg layout estimation via differentiable depth rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12956–12965, June 2021. **2, 3, 4, 5, 6**
- [20] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 8616–8625, 2022. **2**
- [21] Chenggang Yan, Biyao Shao, Hao Zhao, Ruixin Ning, Yongdong Zhang, and Feng Xu. 3d room layout estimation from a single rgb image. *IEEE Transactions on Multimedia*, 22(11):3014–3024, 2020. **2**
- [22] 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 *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019*, pages 3363–3372, 2019. **2, 6**
- [23] Yinda Zhang, Shuran Song, Ping Tan, and Jianxiong Xiao. Panocontext: A whole-room 3d context model for panoramic scene understanding. In *European conference on computer vision*, pages 668–686. Springer, 2014. **1, 2**
- [24] Yining Zhao, Chao Wen, Zhou Xue, and Yue Gao. 3d room layout estimation from a cubemap of panorama image via deep manhattan hough transform. In *European Conference on Computer Vision*, pages 637–654. Springer, 2022. **2**
- [25] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *ECCV*, 2020. **1, 2, 5, 6, 7**
- [26] 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**