

3D Room Layout Estimation from a Cubemap of Panorama Image via Deep Manhattan Hough Transform

Yining Zhao¹, Chao Wen², Zhou Xue², and ✉Yue Gao¹

¹ BNRist, THUICS, BLBCI, KLISS, School of Software, Tsinghua University

² Pico IDL, ByteDance

Abstract. Significant geometric structures can be compactly described by global wireframes in the estimation of 3D room layout from a single panoramic image. Based on this observation, we present an alternative approach to estimate the walls in 3D space by modeling long-range geometric patterns in a learnable Hough Transform block. We transform the image feature from a cubemap tile to the Hough space of a Manhattan world and directly map the feature to the geometric output. The convolutional layers not only learn the local gradient-like line features, but also utilize the global information to successfully predict occluded walls with a simple network structure. Unlike most previous work, the predictions are performed individually on each cubemap tile, and then assembled to get the layout estimation. Experimental results show that we achieve comparable results with recent state-of-the-art in prediction accuracy and performance. Code is available at <https://github.com/Starrah/DMH-Net>.

Keywords: panorama images, room layout estimation, holistic scene structure

1 Introduction

Recovering 3D geometry from a single image is one of the most studied stopics in computer vision. This ill-posed task is generally solved under specific scenarios, with certain assumptions or with prior knowledge. The goal of our work is to reconstruct the 3D room layout from a single panorama image under the Manhattan world assumption.

As described in [4], humans perceive 3D world with great efficiency and robustness by using geometrically salient global structures. For our task, an intuitive way of representing the room layout is to use the wireframe for a typical 3D room, consisting of lines on the horizontal plane and vertical lines denoting the junctions between vertical walls. Although using wireframes to estimate 3D room structure is compact and sparse, it could be challenging for vision algorithms to detect long and thin lines with few appearance clues, especially when lines are heavily occluded in a cluttered room.

✉Corresponding author.




Fig. 1. The processing pipeline of our method. Taking panorama image as input, we first apply Equirectangular-to-Perspective transform to generate cubemap, then utilize Deep Manhattan Hough Transform to predict the positions of the wall-wall, wall-floor and wall-ceiling intersection line in each cubemap tile, and recover 3D room layouts by with post-processing procedures.

Existing methods either model the structure estimation as a segmentation problem in panoramic or perspective [39] or decompose the estimation of the geometry elements into a series of regression problems [33]. The motivation of our work is to take advantage of the compact and efficient representation in the VR area and provide an alternative perspective on this problem. We introduce the cubemap [12] of the panorama image and obtain appropriate visual cues in each cubemap tile. To increase the robustness of the line description, we resort to the Hough transform, which is widely used in line segment detection. It uses two geometric terms, an offset, and an angle, to parameterize lines. As the parameterization of lines is global, the estimation is less prone to be affected by noisy observation and partial occlusions.

In our task, we introduce the Manhattan world assumption to the Hough transform, making the representation even simpler. The wireframe of a room has three types of lines. The first two types are wall-ceiling intersection lines and wall-floor intersection lines, which live on horizontal planes and are perpendicular to each other under the Manhattan world assumption. The third type is vertical lines in the 3D space, which represent the junctions between walls and are perpendicular to the first two types of lines. By taking some pre-processing steps to align the room with the camera coordinate system, the first two types of lines can always be aligned with either the x-axis or the y-axis in the floor/ceiling view.

With the aligned input images, after we adopt the equirectangular to perspective (E2P) transform [12,39] to get the cubemap from the panorama images, it can be proven that a straight line which is along one of the coordinate axes in the camera 3D space is either a horizontal line, a vertical line or a line passing the center of the image in the cubemap tiles. This simplifies the estimation of the wireframe lines greatly as only lines with specific characteristics are needed for detection so that the network can be more concentrated to learn a better line detection model suitable for Manhattan room layout estimation.

In this work, we add the Manhattan world line priors into deep neural networks to overcome the challenge of lacking appearance features and occlusion for 3D room wireframes by relying on Hough transform. We embed the Hough

transform into trainable neural networks in order to combine Manhattan world wireframe priors with local learned appearances.

The main contributions of this paper can be summarized as follows:

- We introduce the Manhattan world assumption through Deep Hough Transform to capture the long-range geometric pattern of room layouts.
- We propose a novel framework estimating layouts on each cubemap tile individually, which is distortion-free for standard CNN.
- We directly predict Manhattan lines with explicit geometric meaning, which achieves comparable performance towards recent state-of-the-art works.

2 Related Work

Room layout estimation. 3D room layout estimation from a single image attracted a lot of research over the past decade. Most previous studies exploit the Manhattan world assumption [7] which means that all boundaries are aligned with the global coordinate system. Moreover, vanishing points detection can be used to inferring the layout based on the assumption.

Traditional methods extract geometric cues and formalize this task as an optimization problem. Since the images may differ in the FoV (field of view), ranging from perspective to 360° panoramas, the methods vary with types of input images. In terms of perspective images, Delage *et al.* [10] propose a dynamic Bayesian network model to recognize “floor-wall” geometry of the indoor scene. Lee *et al.* [23] using Orientation Map (OM) while Hedau *et al.* [16] using Geometric Context (GC) for geometry reasoning to tackle the problem. The strategies have been employed by other approaches leveraging enhanced scoring function [31,30], or modeling the objects-layout interaction [9,13,45].

On the other hand, since the 360° panoramas provides more information, there are multiple papers exploiting in this direction. Zhang *et al.* [43] proposes to estimate layout and 3D objects by combining the OM and GC on a panoramic image. Yang *et al.* [38] takes line segments and superpixel facets as features and iteratively optimize the 3D layout. Xu *et al.* [36] estimate layout use detected objects, their pose, and context in the scene. In order to recover the spatial layout, Yang *et al.* [40] use more geometric cues and semantic cues as input, while Pintore *et al.* [27] utilize the gradient map.

With the astonishing capability, neural network based methods leverage data prior to improve layout estimation. Most studies in this field focused on dense prediction, which train deep classification network to pixel-wise estimate boundary probability map [25,29,44], layout surface class [8,19] or corner keypoints heatmaps [22,11]. The panoramas-based approach has recently attracted wide interest to 3D room layout estimation problems. Zou *et al.* [46] predict the corner and boundary probability map directly from a single panorama. Yang *et al.* [39] propose DuLa-Net which leverages both the equirectangular view and the perspective floor/ceiling view images to produce 2D floor plan. Pintore *et al.* [26] follows the similar approach, which directly adopts equirectangular to perspective (E2P) conversion to infer 2D footprint. Fernandez-Labrador *et*

al. [11] present EquiConvs, a deformable convolution kernel which is specialized for equirectangular image. Zeng *et al.* [42] jointly learns layout prediction and depth estimation from a single indoor panorama image. Wang *et al.* [35] propose a differentiable depth rendering procedure which can learn depth estimation without depth ground truth. Although extensive research has been carried out on 3D room layout estimation, few studies exist which try to produce a more compact representation to infer the layout. Sun *et al.* propose HorizonNet [33] and HoHoNet [34], which encodes the room layout as boundary and corner probability vectors and propose to recover the 3D room layouts from 1D predictions. LGT-Net [20] is a recent work which represents the room layout by horizon-depth and room height. NonCuboidRoom [37] takes lines into account and recovers partial structures from a perspective image, but cannot estimate the whole room. Although these existing methods are sound in predicting layout, they still have limitations in representing layout. Our approach also assembles compact results but is different from the existing methods. Rather than regressing boundaries values, our approach combines the cubemap and parametric representation of lines.

Hough transform based detectors. Hough *et al.* [18] devise Hough transform to detect line or generalized shapes [2] *e.g.* circle from images. Through the extensive use of Hough transform, traditional line detectors apply edge detection filter at the beginning (*e.g.* Canny [5] and Sobel [32]), then identify significant peaks in the parametric spaces. Qi *et al.* [28] using the Hough voting schemes for 3D object detection. Beltrametti *et al.* [3] adopt voting conception for curve detection. Recently, learning-based approaches demonstrate the representation capability of the Hough transform. Han *et al.* [14] propose Hough Transform to aggregate line-wise features and detect lines in the parametric space. Lin *et al.* [24] consider adding geometric line priors through Hough Transform in networks. Our work use Deep Hough features in a new task to tackle the room layout estimation problem.

3 Method

3.1 Overview

Our goal is to estimate the Manhattan room layout from a single 360° panoramic image. However, panoramic images have distortion, *i.e.* a straight line in the 3D space may not be straight in the equirectangular view of panoramic images. Instead of make predictions directly on the panoramic image like previous works [46,33,35,34], we get a cubemap [12] which contains six tiles by adopting E2P transform [12,39].

Given a single RGB panoramic image as input, first we take some pre-processing steps to align the image, get the cubemap and transform ground truth labels. Then we use our proposed Deep Manhattan Hough Network (DMH-Net) to detect three types of straight lines on the tiles of the cubemap. Finally, with an optimization-based post-processing procedure, the line detection results can be assembled and fully optimized to generate the 3D room layout.

In Sec. 3.2, we introduce our pre-processing procedure. The Deep Manhattan Hough Transform for room layout estimation is presented in Sec. 3.3. We summarize the network architecture of the proposed DMH-Net in Sec. 3.4. Finally, in Sec. 3.5, we introduce our optimization-based post-processing method.

3.2 Pre-processing

Aligning the Image. Receiving a single panorama which covers a 360° H-FoV, we first align the image based on the LSD algorithm and vanishing points mentioned in [46,33]. Our approach exploits both the Manhattan world assumption and the properties of aligned panoramas. After the alignment, the cubemap tiles are aligned with the three principal axes [7], *i.e.*, the optical axis of the front camera is perpendicular to a wall. Aligning the panorama makes the vanishing point precisely the center of each cubemap tile.

Cubemap Generation. E2P transform is conducted six times on the equal-rectangular image with different azimuth angle and elevation angle to generate cubemap tiles $I_{front}, I_{back}, I_{left}, I_{right}, I_{ceil}, I_{floor}$, as Fig. 1 shown. For all of the cubemap tiles, FoV is set to 90° on both horizontal and vertical direction, and image size is set to 512×512 .

Ground Truth Transformation. After adopting the aligning method described above for all the datasets we use, the ground truth of the room’s layout is provided in the format of corner coordinates in the panoramic image. Using the same E2P transform as described previously, we can also transform the coordinates of the ground truth corner coordinate from the panoramic image to the cubemap and then get lines in the cubemap by connecting the points by their connectivity relation in the original panoramic image. The lines can be categorized into three types: horizontal lines, vertical lines and lines passing the center of the image. Since Deep Manhattan Hough Transform can only detect straight lines rather than line segments, we do not care about the specific position of the line’s start point and end point. So for a horizontal (or vertical) line, we only use the y (or x) coordinate to represent the line, and for a line passing the center, we only use the orientation angle θ to represent it.

3.3 Deep Manhattan Hough Transform

In the case of line detection, the traditional Hough Transform [18] parameterizes lines in images to polar coordinates with two parameters, an orientation θ and a distance ρ . Every image pixel votes to a discretized parameter bin, which represents the corresponding polar coordinates. The parameter space is denoted as the Hough space, and the maximum local peaks represent the lines in the image. Specifically, given an single channel input $\mathbf{X} \in \mathbb{R}^{h_0 \times w_0}$, the Hough Transform \mathcal{H} can be represented as:

$$\mathcal{H}(\rho, \theta) = \sum_{i \in l} \mathbf{X}(x_i, y_i) \quad (1)$$

in which l is a line whose orientation angle is θ and distance to the coordinate origin is ρ , and i is each of the point in line l .

A key concept of our network is detecting all possible positions of the layout boundaries of the room in each cubemap tile. We propose to incorporate the deep network with Hough Transform for layout boundaries detection. Specifically, we present the Deep Manhattan Hough Transform (DMHT) combining the deep CNN feature and Manhattan world assumption. It is based on the following two assumption:

1. The Manhattan world assumption, *i.e.*, all of the walls, the ceiling and the floor must be perpendicular to each other, and all of the intersection lines of them must be parallel with one of the coordinate axes of some orthogonal coordinate space (named Manhattan space).
2. The input image must be aligned, *i.e.*, the camera of each cubemap tile faces precisely to one of the walls, and its optical axis is parallel with one of the coordinate axes of the Manhattan space.

In practice, the two assumptions are quite straightforward since most of the rooms in human buildings obey the Manhattan world assumption, and the second assumption can be implemented with the pre-processing steps described in Sec. 3.2. Under these two assumptions, it can be proven that any line in the wireframe of the room, including wall-wall line, wall-ceiling line and wall-floor line, must be either a horizontal line ($\theta = 0$), a vertical line ($\theta = \pi/2$) or a line passing the center ($\rho = 0$) in the cubemap tiles, which is a special case of single-point perspective [17] (more details in supplementary materials).




Fig. 2. Deep Manhattan Hough Transform Overview. Each feature map of the CNN can be regarded as many discretized bins. According to the Hough Transform, the bins with the same polar coordinate parameters vote for the same single bin in the Hough space. A bin in the vector **H** (green) is calculated by aggregating horizontal features in the 2D feature map, while bins in **V** (orange) and **C** (blue) are calculated by aggregate vertical features and features passing the center respectively.

As shown in Fig. 2, given the feature maps of a cubemap tile extracted by the encoder network as input, for each channel $\mathbf{X} \in \mathbb{R}^{h \times w}$ of the feature maps, the Deep Manhattan Hough Transform \mathcal{MH} output three vectors **H**, **V** and **C**,

corresponding to bins in the Hough space representing horizontal lines, vertical lines and lines passing the center, which is defined as:

$$\begin{aligned}
 \mathbf{H}(\rho) &= \mathcal{MH}_H(\rho) = \mathcal{H}(\rho, 0) = \sum_{x_i = -\frac{w}{2}}^{\frac{w}{2}} \mathbf{X}(x_i, \rho), \\
 \mathbf{V}(\rho) &= \mathcal{MH}_V(\rho) = \mathcal{H}(\rho, \frac{\pi}{2}) = \sum_{y_i = -\frac{h}{2}}^{\frac{h}{2}} \mathbf{X}(\rho, y_i), \\
 \mathbf{C}(\theta) &= \mathcal{MH}_C(\theta) = \mathcal{H}(0, \theta) \\
 &= \begin{cases} \sum_{x_i=0}^{\frac{w}{2}} \mathbf{X}(x_i, x_i \cdot \tan(\theta)), & |\tan(\theta)| \leq 1 \\ \sum_{y_i=0}^{\frac{h}{2}} \mathbf{X}(y_i \cdot \cot(\theta), y_i), & |\tan(\theta)| > 1 \end{cases}
 \end{aligned} \tag{2}$$

in which $\rho \in [-\frac{h}{2}, \frac{h}{2}]$ for \mathbf{H} and $\rho \in [-\frac{w}{2}, \frac{w}{2}]$ for \mathbf{V} , and $0 \leq \theta \leq 2\pi$. $\mathbf{H}(\rho)$ is the bin of \mathbf{H} with Hough space parameter ρ , and similar for \mathbf{V} and \mathbf{C} .

To calculate the proposed DMHT efficiently, effective discretization is necessary. It is natural to discretize ρ to be integers so that each bin in \mathbf{H} and \mathbf{V} represents to a 1-pixel-wide line in the image, thus $\mathbf{H} \in \mathbb{R}^h$ and $\mathbf{V} \in \mathbb{R}^w$. In our experiment, we discretize θ so that $\mathbf{C} \in \mathbb{R}^{2(h+w)}$ for each bin in \mathbf{C} , the corresponding line intersects the image's border at positions whose coordinates are integer. With the discretization technique above, the process of DMHT can be implemented with matrix addition and multiplications, which is highly parallelizable and suitable for GPU calculation. See the supplementary materials for more detail.

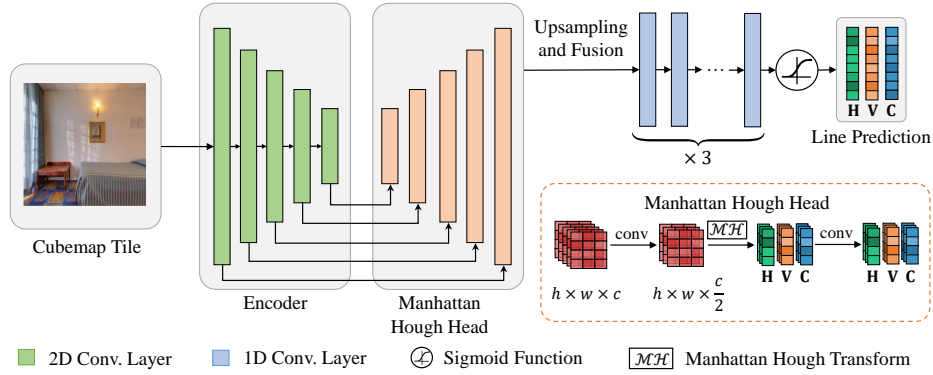


Fig. 3. Overview of the architecture of Deep Manhattan Hough Network(DMH-Net). Given a cubemap tile image as input, a CNN encoder is adopted to extract multi-level image features. Manhattan Hough Head are utilized to handle multi-scale 2D features and perform the Manhattan Hough Transform to get three feature vectors in the Hough space, $\mathbf{H} \in \mathbb{R}^{h \times \frac{c}{2}}$, $\mathbf{V} \in \mathbb{R}^{w \times \frac{c}{2}}$ and $\mathbf{C} \in \mathbb{R}^{2(h+w) \times \frac{c}{2}}$. Finally, the feature vectors are fused to generate lines prediction result, in the format of line position probability vector.

3.4 Network Architecture

Feature Extractor. We employ Dilated Residual Networks [41,15] as our encoder, which utilizes dilated convolution to improve spatial acuity for the images to learn better feature for thin line prediction. The input shape of panorama in equirectangular view I_{equi} is $512 \times 1024 \times 3$. For each of the six tiles of the cubemap, the input shape is $512 \times 512 \times 3$. To capture both low-level and high-level features, we extract the intermediate features of the encoder network. In particular, we gather the perceptual feature from each block before downsample layers. Then, all the features are fed into five independent **Manhattan Hough Head** with different input and output size.

Manhattan Hough Head. As shown in Fig. 3, Manhattan Hough Head receives feature maps of cubemap tile and produces feature vectors in Hough space for line detection. Specifically, given input feature maps with channel number c , the module first applies 2D convolution to reduce the feature channels to $\frac{c}{2}$, then transform the feature through DMHT described in 3.3 to get results in Hough space. The results are then filtered by a 1D convolution layer with kernel size = 3 to capture context information surrounding certain Hough bins. The output of this module is three multi-channel feature vectors in the Hough space: $\mathbf{H} \in \mathbb{R}^{h \times \frac{c}{2}}$, $\mathbf{V} \in \mathbb{R}^{w \times \frac{c}{2}}$ and $\mathbf{C} \in \mathbb{R}^{2(h+w) \times \frac{c}{2}}$.

Upsampling, Fusion and Generating Line Prediction. Since the size of the feature maps extracted by the feature extractor varies with the depth of the layer, the size of the output feature vectors of the five Manhattan Hough Heads differs. We upsample all the feature vectors to the same size which equals to the width and height of the original image, $h = 512$ and $w = 512$, with bilinear interpolation. Then, feature vectors of the same type are fused by channel-wise concatenation, and filtered by three 1D convolution layers. Finally, the Sigmoid function is applied to generate three single-channel 1D vectors $\mathbf{H} \in \mathbb{R}^h$, $\mathbf{V} \in \mathbb{R}^w$ and $\mathbf{C} \in \mathbb{R}^{2(h+w)}$, which represents the predicted probability of horizontal lines, vertical lines and lines passing the center of the image.

Loss Functions. If we define the ground truth of the probability vectors \mathbf{H} , \mathbf{V} and \mathbf{C} as binary-valued vectors with 0/1 labels simply, it would be too sparse to train, *e.g.* there would be only less than two non-zero values out of 512 in a cubemap tile. Similar as HorizonNet[33], we smooth the ground truth based on the exponential function of the distance to nearest ground truth line position.

Then, for \mathbf{H} , \mathbf{V} and \mathbf{C} , we can apply binary cross entropy loss:

$$\mathcal{L}_{bce}(\mathbf{X}, \mathbf{X}^*) = - \sum_i x_i^* \log(x_i) + (1 - x_i^*) \log(1 - x_i) \quad (3)$$

in which x_i denotes the i -th element of \mathbf{X} .

The overall loss for layout estimation is defined by the sum of the loss of the three type of lines:

$$\mathcal{L} = \mathcal{L}_{bce}(\mathbf{H}, \mathbf{H}^*) + \mathcal{L}_{bce}(\mathbf{V}, \mathbf{V}^*) + \mathcal{L}_{bce}(\mathbf{C}, \mathbf{C}^*) \quad (4)$$

where the notation $*$ indicates ground truth data.

3.5 Post-processing

Since the output of our network is line existence confidence on the cubemap, post-processing is necessary to generate the final 3D layout result. Our post-processing procedure consists of two stages: initialization and optimization (more details in supplementary materials).

Parametric Room Layout Representation. In 3D space, take the camera as the coordinate origin and z-axis points to the ceiling, and regard the distance from the camera to the floor can be set as a known value similar to Zou *et al.* [46]. The layout of a Manhattan room with n vertical walls can be represents with $n + 1$ parameters: n for the distance on coordinate axis from the camera to each wall, the other for the height of the room.

Layout Parameters Initialization. Given the lines predictions in the cube-map, we can calculate the elevation angle of each wall-ceiling and wall-floor intersection line, and since the distance from the camera to the floor is known, the distance to each of the wall-floor lines can be directly estimated, which is used as initial values for the first n parameters of the layout representation. Though the distance from the camera to the ceiling is unknown, by assuming and optimizing the ratio of the distance from the camera to the ceiling and the floor so that the ceiling 2D frame has the highest 2DIoU with the floor 2D frame, we can get the initial estimated value of the room height.

Gradient Descent Based Optimization. Given the initialized layout parameters, we convert the them into corner point coordinates in the panoramic image and then line positions in the cubemap. For each line position in the cubemap tiles, we can get line existence confidence from the output H , V and C of the network. We add the confidence of all the line position together to get a score representing the overall confidence of the layout parameters, define the loss as the negative of the score and optimize the loss using SGD, with learning rate 0.01, no momentum and 100 optimization steps. With the optimization process, the line predictions on different tiles of the cubemap can be effectively integrated together to get a better layout estimation from the global aspect.

4 Experiments

4.1 Experimental setup

Dataset. We use three different datasets. The PanoContext dataset [43] and the Stanford 2D-3D dataset [1,46] are used for cuboid layout estimation, while the Matterport 3D dataset [6] is used for non-cuboid Manhattan layout estimation. For fair comparison, we follow the train/validation/test split adopted by other works [46,33,47].

Evaluation Metric We employ standard evaluation metrics for 3D room layout estimation. For both cuboid and non-cuboid estimation, we use (i) **3D IoU**, which is defined as the volumetric intersection over the union between the prediction and the ground truth. In addition, for cuboid layout estimation, we also use (ii) **Corner Error** (CE), which measures the average Euclidean distance

| Method | PanoContext | | | Stanford 2D-3D | | |
|----------------------------|--------------|-------------|-------------|----------------|-------------|-------------|
| | 3D IoU | CE | PE | 3D IoU | CE | PE |
| PanoContext [43] | 67.23 | 1.60 | 4.55 | - | - | - |
| LayoutNet [46] | 74.48 | 1.06 | 3.34 | 76.33 | 1.04 | 2.70 |
| DuLa-Net [39] | 77.42 | - | - | 79.36 | - | - |
| CFL [11] | 78.79 | 0.79 | 2.49 | - | - | - |
| HorizonNet [33] | 82.17 | 0.76 | 2.20 | 79.79 | 0.71 | 2.39 |
| AtlantaNet [11] | - | - | - | 82.43 | 0.70 | 2.25 |
| LED ² -Net [35] | 82.75 | - | - | 83.77 | - | - |
| DMH-Net (Ours) | 85.48 | 0.73 | 1.96 | 84.93 | 0.67 | 1.93 |

Table 1. Quantitative results of cuboid room layout estimation evaluated on the testset of the PanoContext dataset [43] and the Stanford 2D-3D dataset [1,46]. CE means Corner Error, and PE means Pixel Error.

between predicted corners and ground truth, and **(iii) Pixel Error (PE)**, which is the pixel-wise accuracy. For non-cuboid estimation, we also use **(iv) 2D IoU**, which project 3D IoU to the 2D plane, and **(v) δ_i** , which is defined as the percentage of pixels where the ratio between the prediction and the ground truth is no more than 1.25. For 3D IoU, 2D IoU and δ_i , the larger is better. For Corner Error and Pixel Error, the smaller is better.

Baselines We compare to previous works for both cuboid and non-cuboid room layout estimation and evaluate quantitative performance. The methods we compare to include PanoContext [43], LayoutNet [46], CFL [11], DuLa-Net [39], HorizonNet [33], AtlantaNet [26], HoHoNet[34] and LED²-Net [35]. In addition, we also compared the improved version of the algorithm proposed in the recent work of Zou *et al.* [47], *i.e.*, Layoutnet-v2, DuLa-Net-v2 and HorizonNet⁺. We compare with all these methods based on published comparable results. More specifically, DuLa-Net [39] and CFL [11] take 256×512 resolution images as input while others using the input resolution of 512×1024 . Additionally, PanoContext [43] and LayoutNet [46] need additional input (*i.e.* Line Segment or Orientation Map) besides a single RGB panorama image.

Implementation Details The network is implemented in PyTorch and optimized using Adam [21]. The network is trained for 75 epochs with batch size 8 and learning rate $1e-4$. The model is trained on a single NVIDIA A100 GPU. For both cuboid and non-cuboid estimation, we use the Pano Stretch[33], left-right flipping for data augmentation. In addition, for non-cuboid estimation, we also apply vertical flip and random 90° rotation augmentation on the cubemap tiles generated by the E2P transform.

4.2 Cuboid Room Results

Tab. 1 show the quantitative comparison on the PanoContext dataset [43] and Stanford 2D-3D dataset [1,46], respectively. Blank fields in the tables indicate metrics which is not reported in the corresponding publications. As can be seen, our approach achieves the state-of-the-art results under these two datasets in

cuboid 3D room layout estimation. The qualitative results of the PanoContext and the Stanford 2D-3D dataset are shown in Fig. 4.

4.3 Non-cuboid Room Results

Tab. 2 shows the quantitative comparison on the Matterport 3D dataset [6]. Blank fields in the table indicate metrics which is not reported in the corresponding publications. As can be seen, our approach achieves competitive overall performance with the state-of-the-art methods in non-cuboid 3D layout estimation. The qualitative results of the Matterport 3D dataset are shown in Fig. 4. We present results from different methods aligned with ground truth. Please see supplemental materials for more qualitative results.

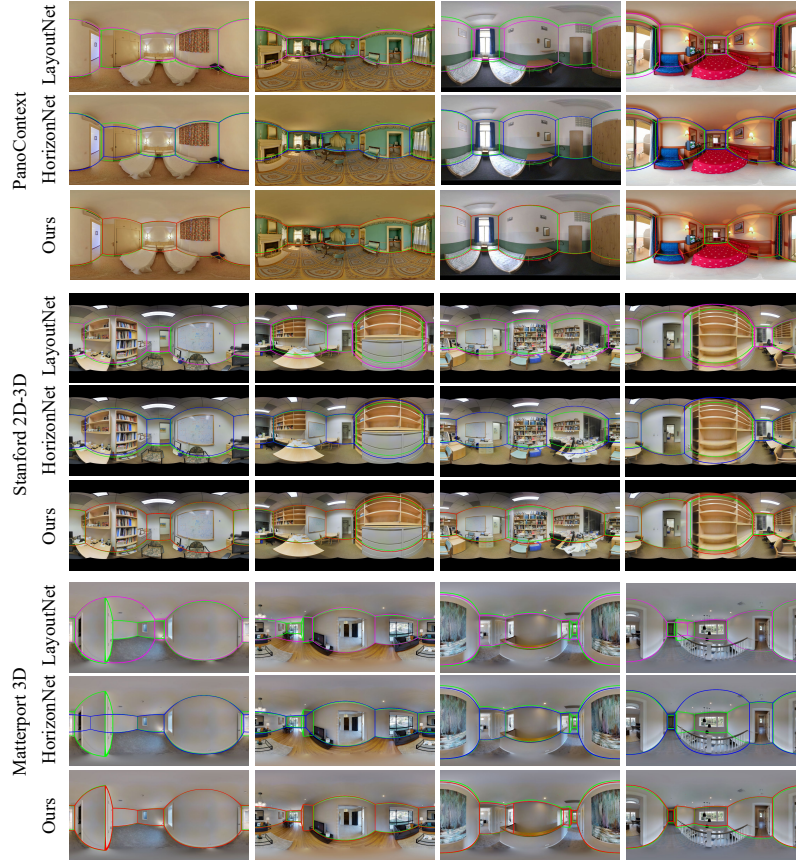


Fig. 4. Qualitative results of both cuboid and non-cuboid layout estimation. The results are selected from four folds that comprise results with the best 0–25%, 25–50%, 50–75% and 75–100% 3DIoU (displayed from left to right). The green lines are ground truth layout while the pink, blue and orange lines are estimated by LayoutNet [46], HorizonNet [33] and our DMH-Net.

| Metrics | 3D IoU | | | | | 2D IoU | | | | | δ_i | | | | | |
|---------------------------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|--------------|
| | # of corners | 4 | 6 | 8 | 10+ | Overall | 4 | 6 | 8 | 10+ | Overall | 4 | 6 | 8 | 10+ | Overall |
| LayoutNet-v2 [47] | | 81.35 | 72.33 | 67.45 | 63 | 75.82 | 84.61 | 75.02 | 69.79 | 65.14 | 78.73 | 0.897 | 0.827 | 0.877 | 0.8 | 0.871 |
| DuLa-Net-v2 [47,39] | | 77.02 | 78.79 | 71.03 | 63.27 | 75.05 | 81.12 | 82.69 | 74 | 66.12 | 78.82 | 0.818 | 0.859 | 0.823 | 0.741 | 0.818 |
| HorizonNet ⁺ [47,33] | | 81.88 | 82.26 | 71.78 | 68.32 | 79.11 | 84.67 | 84.82 | 73.91 | 70.58 | 81.71 | 0.945 | 0.938 | 0.903 | 0.861 | 0.929 |
| AtlantaNet [26] | | 82.64 | 80.1 | 71.79 | 73.89 | 81.59 | 85.12 | 82.00 | 74.15 | 76.93 | 84.00 | 0.950 | 0.815 | 0.911 | 0.915 | 0.945 |
| HoHoNet [34] | | 82.64 | 82.16 | 73.65 | 69.26 | 79.88 | 85.26 | 84.81 | 75.59 | 70.98 | 82.32 | - | - | - | - | - |
| LED ² -Net [35] | | 84.22 | 83.22 | 76.89 | 70.09 | 81.52 | 86.91 | 85.53 | 78.72 | 71.79 | 83.91 | - | - | - | - | - |
| DMH-Net (Ours) | | 84.39 | 80.22 | 66.15 | 64.46 | 78.97 | 86.94 | 82.31 | 67.99 | 66.2 | 81.25 | 0.949 | 0.951 | 0.838 | 0.864 | 0.925 |

Table 2. Quantitative results of non-cuboid room layout estimation evaluated on the Matterport 3D dataset [6].

4.4 Ablation Study

In this section, we perform an ablation study to demonstrate the effect of the major designs in our method.

Choice of Manhattan Line Types We first study the necessity of the three types of lines in the DMHT. We conduct experiments with each of the types of line detection omitted, respectively. The results are summarized in Tab 3. All of the types of line detection are important as disabling any of each results in an obvious performance loss. Among them, horizontal lines detection is more important, as horizontal lines make up the ceiling and floor wireframe which is essential to the initialization in the post-processing procedure.

Cubemap and DMHT We then verify whether the cubemap representation and DMHT helps to recover 3D room layout. We analyze the effectiveness of our full pipeline (*i.e.* III) with another two ablation variants in Tab. 4. The variation (I) encodes feature from cubemap as ours but replaces decoder from DMHT to HorizonNet-like LSTM [33]. The significant performance drop without DMHT in (I) indicates the necessity of DMHT in our full pipeline. We also compare the deep Hough feature against classical Hough Transforms line detectors, please refer to supplementary materials for more details. The variation (II) directly extracts features from the panorama image then apply E2P transform to feed features into DMHT. It shows the effectiveness of cubemap representation since the distortion in equirectangular features can degrade performance. These ablations between (I, II) and (III) provide evidence that the performance gain owe to the use of cubemap and DMHT.

| Line Detection | | | 3D IoU | CE | PE |
|----------------|----------|----------|--------------|-------------|-------------|
| H | V | C | | | |
| × | ✓ | ✓ | 43.58 | 3.69 | 11.91 |
| ✓ | × | ✓ | 83.28 | 0.89 | 2.44 |
| ✓ | ✓ | × | 84.15 | 0.79 | 1.96 |
| ✓ | ✓ | ✓ | 85.48 | 0.73 | 1.96 |

Table 3. Ablation study on the impact of different types of line detection on the PanoContext dataset, in which **H** means horizontal line detection, **V** means vertical line detection and **C** means detection of lines passing the center.

4.5 Model Analysis

Encoder Network We first investigate how our cubemap representation and DMHT are model agnostic. We choose four different encoder backbones: ResNet-18,34,50[15] and DRN-38[41]. We change backbone for both our network and HorizonNet [33] baseline, and then train and evaluate the networks on the PanoContext dataset. As shown in Tab 5, our method is not only effective for the DRN-38 architecture we use, but also consistently outperforms the competitors quantitatively. Ours with Res-50 is even slightly higher than DRN, may owing to deeper layers.

Robustness to occlusion We then test how our method outperforms competitors w.r.t. occlusion scenarios. We provide qualitative comparison in Fig. 5(a). As illustrated, our method has more accurate prediction results in regions where walls are occluded by other clutter. We further set up a more challenging experiment. We manually add challenging noise patches, and then test the network without re-training. As shown in Fig. 5(b), the raw output of ours is more reasonable than competitors in noisy regions. Moreover, quantitative results on the PanoContext dataset also show that our method achieves a 3D IoU of 84.39 while HorizonNet is 78.74.

Limitations We show two examples for analyzing our performance limitations on non-cuboid data. Fig. 6(a) shows that our method may suffer from discriminating two very close walls. Fig. 6(b) shows low line confidence due to “Vote Splitting”. The probability peaks of short lines (in leftmost tile) are less significant than long ones. The short ones may not have enough feature bins to vote

| Method | encoder | | decoder | | 3D IoU |
|------------|---------|-----------------|---------|-----|--------------|
| | cubemap | equirectangular | DMHT | RNN | |
| I | ✓ | | | ✓ | 81.44 |
| II | | ✓ | ✓ | | 83.99 |
| III (Ours) | ✓ | | ✓ | | 85.48 |

Table 4. Ablation study on encoder feature representation and decoder choice. We analyze the effectiveness of our pipeline (*i.e.* III) with another two ablation variants in Tab. 4. Our pipeline (*i.e.* III) achieves best quantitative results.




Fig. 5. Qualitative comparison in occlusion scenarios. (a) shows the result after post-processing, (b) shows the raw output of the network with challenging noise.

| Method | Ours | | | HorizonNet [33] | | |
|-----------|--------------|-------------|-------------|-----------------|------|------|
| Backbone | 3DIoU | CE | PE | 3DIoU | CE | PE |
| ResNet-18 | 82.96 | 0.83 | 2.18 | 79.02 | 1.52 | 2.77 |
| ResNet-34 | 83.84 | 0.87 | 2.13 | 79.97 | 0.91 | 2.49 |
| ResNet-50 | 85.90 | 0.64 | 1.75 | 82.17 | 0.76 | 2.20 |
| DRN-38 | 85.48 | 0.73 | 1.96 | 81.52 | 0.79 | 2.19 |

Table 5. Model analysis of the model agnostic generalization capability of DMH-Net on the PanoContext dataset. (Bold and underline indicates better than other competitors and the best among all, respectively.) Our DMH-Net consistently outperforms the competitors quantitatively when changing different backbones.

for. These cases may be improved by exploring adding interactions between tiles and an improved version of DMHT, which will be our future works.




Fig. 6. Non-cuboid results explanations.

5 Conclusion

In this paper, we introduced a new method for estimating 3D room layouts from a cubemap of panorama image by leveraging DMHT. The proposed method detects horizontal lines, vertical lines and lines passing the image center on the cubemap and applies post-processing steps to combine the line prediction results to get the room layout. The learnable Deep Hough transform enables the network to capture the long-range geometric pattern and precisely detect the lines in the wireframe of the room layout. Both quantitative and qualitative results demonstrate that our method achieves better results for cuboid room estimation and comparable results for non-cuboid room estimation with recent state-of-the-art methods in prediction accuracy and performance. For future work, combining with fully differentiable post-processing or Atlanta world assumption are some of the practical directions to explore.

Acknowledgements This work was supported by National Natural Science Funds of China (No. 62088102, 62021002) and ByteDance Research Collaboration Project.

References

1. Armeni, I., Sax, S., Zamir, A.R., Savarese, S.: Joint 2d-3d-semantic data for indoor scene understanding. arXiv preprint arXiv:1702.01105 (2017) [9](#), [10](#)
2. Ballard, D.H.: Generalizing the hough transform to detect arbitrary shapes. *Pattern recognition* **13**(2), 111–122 (1981) [4](#)
3. Beltrametti, M.C., Campi, C., Massone, A.M., Torrente, M.: Geometry of the hough transforms with applications to synthetic data. *Mathematics in Computer Science* pp. 1–23 (2020) [4](#)
4. Bertamini, M., Helmy, M., Bates, D.: The visual system prioritizes locations near corners of surfaces (not just locations near a corner). *Attention, Perception, & Psychophysics* **75**(8), 1748–1760 (2013) [1](#)
5. Canny, J.: A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8**(6), 679–698 (1986) [4](#)
6. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3D: Learning from RGB-D data in indoor environments. *Proc. of the International Conf. on 3D Vision (3DV)* pp. 667–676 (2017) [9](#), [11](#), [12](#)
7. Coughlan, J.M., Yuille, A.L.: Manhattan world: Compass direction from a single image by bayesian inference. In: *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*. pp. 941–947 (1999) [3](#), [5](#)
8. Dasgupta, S., Fang, K., Chen, K., Savarese, S.: Delay: Robust spatial layout estimation for cluttered indoor scenes. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. pp. 616–624 (2016) [3](#)
9. Del Pero, L., Bowdish, J., Kermgard, B., Hartley, E., Barnard, K.: Understanding bayesian rooms using composite 3d object models. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. pp. 153–160 (2013) [3](#)
10. Delage, E., Lee, H., Ng, A.Y.: A dynamic bayesian network model for autonomous 3d reconstruction from a single indoor image. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. pp. 2418–2428 (2006) [3](#)
11. Fernandez-Labrador, C., Facil, J.M., Perez-Yus, A., Demonceaux, C., Civera, J., Guerrero, J.J.: Corners for layout: End-to-end layout recovery from 360 images. *IEEE Robotics and Automation Letters (RA-L)* **5**(2), 1255–1262 (2020) [3](#), [4](#), [10](#)
12. Greene, N.: Environment mapping and other applications of world projections. *IEEE computer graphics and Applications* **6**(11), 21–29 (1986) [2](#), [4](#)
13. Gupta, A., Hebert, M., Kanade, T., Blei, D.: Estimating spatial layout of rooms using volumetric reasoning about objects and surfaces. *Advances in Neural Information Processing Systems (NeurIPS)* pp. 1288–1296 (2010) [3](#)
14. Han, Q., Zhao, K., Xu, J., Cheng, M.M.: Deep hough transform for semantic line detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 249–265 (2021) [4](#)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016) [8](#), [13](#)
16. Hedau, V., Hoiem, D., Forsyth, D.: Recovering the spatial layout of cluttered rooms. In: *Proc. of the IEEE International Conf. on Computer Vision (ICCV)*. pp. 1849–1856 (2009) [3](#)
17. Horry, Y., Anjo, K.I., Arai, K.: Tour into the picture: using a spidery mesh interface to make animation from a single image. In: *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. pp. 225–232 (1997) [6](#)

18. Hough, P.V.: Method and means for recognizing complex patterns (Dec 18 1962), uS Patent 3,069,654 [4](#), [5](#)
19. Izadinia, H., Shan, Q., Seitz, S.M.: Im2cad. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 2422–2431 (2017) [3](#)
20. Jiang, Z., Xiang, Z., Xu, J., Zhao, M.: Lgt-net: Indoor panoramic room layout estimation with geometry-aware transformer network. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 1654–1663 (2022) [4](#)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: Proc. of the International Conf. on Learning Representations (ICLR) (2015) [10](#)
22. Lee, C.Y., Badrinarayanan, V., Malisiewicz, T., Rabinovich, A.: Roomnet: End-to-end room layout estimation. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV). pp. 4875–4884 (2017) [3](#)
23. Lee, D.C., Hebert, M., Kanade, T.: Geometric reasoning for single image structure recovery. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 2136–2143 (2009) [3](#)
24. Lin, Y., Pintea, S.L., van Gemert, J.C.: Deep hough-transform line priors. In: Proc. of the European Conf. on Computer Vision (ECCV). pp. 323–340 (2020) [4](#)
25. Mallya, A., Lazebnik, S.: Learning informative edge maps for indoor scene layout prediction. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV). pp. 936–944 (2015) [3](#)
26. Pintore, G., Agus, M., Gobbetti, E.: Atlantonet: Inferring the 3d indoor layout from a single 360° image beyond the manhattan world assumption. In: Proc. of the European Conf. on Computer Vision (ECCV). pp. 432–448 (2020) [3](#), [10](#), [12](#)
27. Pintore, G., Garro, V., Ganovelli, F., Gobbetti, E., Agus, M.: Omnidirectional image capture on mobile devices for fast automatic generation of 2.5 d indoor maps. In: Proc. IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 1–9 (2016) [3](#)
28. Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: Proc. of the IEEE International Conf. on Computer Vision (ICCV). pp. 9276–9285 (2019) [4](#)
29. Ren, Y., Li, S., Chen, C., Kuo, C.C.J.: A coarse-to-fine indoor layout estimation (cfile) method. In: Asia Conference on Computer Vision (ACCV). pp. 36–51 (2016) [3](#)
30. Schwing, A.G., Hazan, T., Pollefeys, M., Urtasun, R.: Efficient structured prediction for 3d indoor scene understanding. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 2815–2822 (2012) [3](#)
31. Schwing, A.G., Urtasun, R.: Efficient exact inference for 3d indoor scene understanding. In: Proc. of the European Conf. on Computer Vision (ECCV). pp. 299–313 (2012) [3](#)
32. Sobel, I., Feldman, G.: A 3x3 isotropic gradient operator for image processing. a talk at the Stanford Artificial Project in pp. 271–272 (1968) [4](#)
33. Sun, C., Hsiao, C.W., Sun, M., Chen, H.T.: Horizonnet: Learning room layout with 1d representation and pano stretch data augmentation. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 1047–1056 (2019) [2](#), [4](#), [5](#), [8](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
34. Sun, C., Sun, M., Chen, H.T.: Hohonet: 360 indoor holistic understanding with latent horizontal features. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 2573–2582 (2021) [4](#), [10](#), [12](#)
35. Wang, F.E., Yeh, Y.H., Sun, M., Chiu, W.C., Tsai, Y.H.: Led2-net: Monocular 360deg layout estimation via differentiable depth rendering. In: Proc. IEEE Conf.

- on Computer Vision and Pattern Recognition (CVPR). pp. 12956–12965 (2021) 4, 10, 12
36. Xu, J., Stenger, B., Kerola, T., Tung, T.: Pano2cad: Room layout from a single panorama image. In: Proc. IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 354–362 (2017) 3
 37. Yang, C., Zheng, J., Dai, X., Tang, R., Ma, Y., Yuan, X.: Learning to reconstruct 3d non-cuboid room layout from a single rgb image. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2534–2543 (2022) 4
 38. Yang, H., Zhang, H.: Efficient 3d room shape recovery from a single panorama. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 5422–5430 (2016) 3
 39. Yang, S.T., Wang, F.E., Peng, C.H., Wonka, P., Sun, M., Chu, H.K.: Dula-net: A dual-projection network for estimating room layouts from a single rgb panorama. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 3363–3372 (2019) 2, 3, 4, 10, 12
 40. Yang, Y., Jin, S., Liu, R., Bing Kang, S., Yu, J.: Automatic 3d indoor scene modeling from single panorama. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 3926–3934 (2018) 3
 41. Yu, F., Koltun, V., Funkhouser, T.: Dilated residual networks. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 636–644 (2017) 8, 13
 42. Zeng, W., Karaoglu, S., Gevers, T.: Joint 3d layout and depth prediction from a single indoor panorama image. In: Proc. of the European Conf. on Computer Vision (ECCV). pp. 666–682 (2020) 4
 43. Zhang, Y., Song, S., Tan, P., Xiao, J.: Panocontext: A whole-room 3d context model for panoramic scene understanding. In: Proc. of the European Conf. on Computer Vision (ECCV). pp. 668–686 (2014) 3, 9, 10
 44. Zhao, H., Lu, M., Yao, A., Guo, Y., Chen, Y., Zhang, L.: Physics inspired optimization on semantic transfer features: An alternative method for room layout estimation. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 870–878 (2017) 3
 45. Zhao, Y., Zhu, S.C.: Scene parsing by integrating function, geometry and appearance models. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 3119–3126 (2013) 3
 46. Zou, C., Colburn, A., Shan, Q., Hoiem, D.: Layoutnet: Reconstructing the 3d room layout from a single rgb image. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR). pp. 2051–2059 (2018) 3, 4, 5, 9, 10, 11
 47. Zou, C., Su, J.W., Peng, C.H., Colburn, A., Shan, Q., Wonka, P., Chu, H.K., Hoiem, D.: Manhattan room layout reconstruction from a single 360° image: A comparative study of state-of-the-art methods. *International Journal of Computer Vision (IJCV)* **129**(5), 1410–1431 (2021) 9, 10, 12