PolyRoom: Room-aware Transformer for Floorplan Reconstruction

Yuzhou Liu^{1,2,3}[©], Lingjie Zhu⁴, Xiaodong Ma⁵, Hanqiao Ye^{1,2,3}[©], Xiang Gao^{1,3}[⊠] [©], Xianwei Zheng⁶[®], and Shuhan Shen^{1,2,3}[⊠] [©]

 ¹ Institute of Automation, Chinese Academy of Sciences
 ² School of Artificial Intelligence, University of Chinese Academy of Sciences
 ³ CASIA-SenseTime Research Group ⁴ Cenozoic Robotics
 ⁵ Ocean University of China ⁶ The State Key Lab. LIESMARS, Wuhan University {liuyuzhou2021, yehanqiao2022, xiang.gao}@ia.ac.cn; lingjie.zhu.me@gmail.com; maxiaodong5000@stu.ouc.edu.cn; zhengxw27@foxmail.com; shshen@nlpr.ia.ac.cn



(a) Point cloud (b) Density map (c) Initial shapes (d) Refined shapes (e) Floorplan

Fig. 1: Structured floorplan reconstruction by PolyRoom. Given an indoor point cloud (a), PolyRoom predicts the initial room shapes (c) from the density map (b). Then, PolyRoom refines them gradually (d) and reconstructs the final vectorized floorplan (e).

Abstract. Reconstructing geometry and topology structures from raw unstructured data has always been an important research topic in indoor mapping research. In this paper, we aim to reconstruct the floorplan with a vectorized representation from point clouds. Despite significant advancements achieved in recent years, current methods still encounter several challenges, such as missing corners or edges, inaccuracies in corner positions or angles, self-intersecting or overlapping polygons, and potentially implausible topology. To tackle these challenges, we present PolyRoom, a room-aware Transformer that leverages uniform sampling representation, room-aware query initialization, and room-aware self-attention for floorplan reconstruction. Specifically, we adopt a uniform sampling floorplan representation to enable dense supervision during training and effective utilization of angle information. Additionally, we propose a room-aware query initialization scheme to prevent non-polygonal sequences and introduce room-aware self-attention to

 $^{\[}top]$ Corresponding authors

enhance memory efficiency and model performance. Experimental results on two widely used datasets demonstrate that PolyRoom surpasses current state-of-the-art methods both quantitatively and qualitatively. Our code is available at: https://github.com/3dv-casia/PolyRoom/.

1 Introduction

Structured modeling [21, 40] aims to represent a scene using basic geometric primitives while maintaining the complete topology and precise geometry holistically. Different from other commonly employed representations such as point clouds, meshes, or voxels, structured representations are lightweight and editable, facilitating human scene understanding and downstream applications, such as VR/AR, robotics, and navigation [5, 26, 29]. However, efficiently reconstructing structures from raw unstructured data remains an open problem.

For structured indoor modeling, various representations are utilized, ranging from 2D layouts [25] and 3D wireframes [35] to 2D floorplans [7], aiming to reconstruct scenes from diverse data sources such as single monocular RGB images [20], panoramic images [22], CAD drawings [42], and point clouds [3]. Among different input data and output representations, vectorized floorplan exhibit significant advantages in accurately depicting scene geometric details and global topology structure, while point clouds are deemed to be more geometrically precise and contain rich visual information. Consequently, the trend of using point clouds as input and predicting vectorized 2D floorplan has gained popularity in indoor mapping [7, 39]. However, in the face of noise, incompleteness, and other inherent challenges in raw point clouds, floorplan reconstruction remains a demanding task.

To address the challenges of reconstructing structured floorplan from 3D point clouds, there has been a steady paradigm shift from traditional geometric optimization methods to end-to-end inference approaches. Traditional optimization-based methods involve fitting planes from point clouds and obtaining the final floorplan through optimization techniques [3, 15]. To simplify the problem, it is a common practice to project the 3D indoor point clouds along the gravity axis and proceed with the resulting 2D density map. Some approaches propose combining deep learning and geometric optimization techniques by initially detecting room proposals and then solving it as an optimization problem [7,33], which are not end-to-end and rely on prior-based optimization method design. Recently, with Transformer [36] showcasing remarkable capabilities in sequence modeling [31], researchers have attempted to adopt it for indoor mapping tasks, and it also yields promising results [8,39].

For the Transformer-based methods, the floorplan representation has critical impacts on the choice of techniques and the organization of the algorithm. HEAT [8] represents floorplan as discrete corners and edges. However, this representation suffers from missing corners and edges due to missed detection. SLIBO-Net [34] adopts a slicing box representation but is limited by the Manhattan assumption. Representing floorplan as polygon sequences eliminates reliance on the Manhattan assumption. RoomFormer [39] first introduces the two-level queries, which depict the floorplan as variable-size polygons represented by variablelength ordered corners. Despite significant improvements in reconstruction results, RoomFormer [39] still faces challenges such as disordered sequences from random query initialization, as well as inaccuracies in polygons due to missing or biased single corners. Overall, current methods still encounter challenges with incomplete and inaccurate reconstruction results.

In this paper, we propose a novel floorplan reconstruction method named PolyRoom, designed to address floorplan reconstruction as a refinement process of room queries from initial shapes (as depicted in Figure 1). PolyRoom introduces the uniform sampling floorplan representation, enabling room-aware query initialization with fixed-length vertex sequences. Furthermore, it incorporates dense supervision during training and utilizes angles to enhance performance. To mitigate inaccurate polygon sequences resulting from randomly initialized learnable queries, PolyRoom introduces room-aware query initialization to align the queries with room semantics. With the initialization, the initial shapes of room polygons are refined layer by layer in the Transformer decoder. Considering the memory overhead of self-attention, PolyRoom utilizes room-aware self-attention to achieve better performance with reduced memory consumption.

We evaluate PolyRoom on two widely adopted datasets, Structured3D [41] and SceneCAD [1]. On both datasets, PolyRoom surpasses the state of the arts. Besides, PolyRoom also shows better generalization ability. We outline our contributions below:

- We propose PolyRoom, a novel method that views floorplan reconstruction as a progressive refinement of initial room shapes in an encoder-decoder architecture.
- The uniform sampling representation depicts room polygons as fixed-length vertex sequences, providing various dense supervision during training and facilitating the introduction of room-aware query initialization.
- Room-aware query initialization utilizing instance segmentation yields superior initial queries compared with randomly initialized learnable queries, enhancing the accuracy of predicted polygon sequences.
- Room-aware self-attention incorporates intra-room and inter-room self-attention based on the floorplan representation, reducing memory consumption while achieving better reconstruction results.

2 Related Work

2.1 Vectorized Floorplan Reconstruction

Floorplan reconstruction entails generating a structured representation of indoor scenes from unstructured data. Various approaches have been explored in this field. Some researchers fit planes from point clouds and select lines to form the floorplan using optimization techniques [3, 15]. Others infer floorplan from panoramic images [22, 35] or CAD drawings [13, 42] in end-to-end manners.

In recent years, there has been a growing interest in floorplan reconstruction with density maps generated by point cloud projection. Some methods integrate instance segmentation with optimization algorithms to address this challenge [7,33]. However, they are constrained by the prior-based optimization function design and long optimization time. Given the remarkable success of applying Transformer [36] in computer vision [11,31], floorplan reconstruction has also embraced its adoption. HEAT [8] uses a Transformer to detect corners and infer connection relationships between them to create floorplan planar graph, but it may fail when certain corners are absent, leading to missing edges and unclosed rooms. RoomFormer [39] represents floorplan as two-level queries, predicting the polygon sequences to ensure the room closeness. Nonetheless, it faces challenges such as incorrect polygon sequences and distorted polygon contours. Recently, PolyDiffuse [6] first uses diffusion models to generate refined floorplan based on initial proposals or human annotations. SLIBO-Net [34] employs the slicing box representation and introduces a regularization mechanism and post-processing step with geometric priors to capture finer local geometric details.

2.2 Polygonal Instance Segmentation

Floorplan reconstruction using polygon sequence representation treats rooms as polygons, akin to polygonal instance segmentation which predicts polygonal instance contours. Some researchers extract polygon instances based on geometric primitives. For instance, Delaunay point processes [14] embed point processes into a Delaunay triangulation and utilize its key properties in a Markov Chain Monte Carlo sampler [2] to extract geometric structures from images. Besides the optimization-based methods, an increasing number of end-to-end polygonal instance segmentation algorithms have emerged, which typically employ timeseries models to predict polygon boundaries. Castrejon et al. [4] uses LSTM [18] to sequentially produce polygon vertices of object contours. PolyTransform [27] employs a segmentation network to generate instance masks and refines initial contours with a deforming network to better fit the object contours. Deep-Snake [32] implements the snake algorithm by iteratively deforming an initial contour to match the object boundary. BoundaryFormer [24] utilizes a Transformer architecture to directly predict polygon contours using a mask-based loss through a differentiable rasterizer. Recently, PolyFormer [30] predicts the sequence of polygon vertices with large language models by combining image patches and text query tokens.

3 Method

The proposed PolyRoom views the floorplan reconstruction problem as a progressive refinement process of room queries, which is shown in Figure 2. In this section, we begin by introducing our floorplan representation, which is followed by the description of the encoder and decoder modules of PolyRoom. Subsequently, we present the room query initialization module and the floorplan extraction module. Finally, we detail the loss functions employed in PolyRoom.



Fig. 2: Overall architecture of PolyRoom. PolyRoom consists of four main components: (a) Encoder module, (b) Decoder module, (c) Room-aware query initialization module, and (d) Floorplan extraction module. Room queries are initialized with instance segmentation. Subsequently, they are refined in the Transformer decoder layer by layer with dense supervision (red and blue boxes mark the changes). Finally, the floorplan is extracted based on vertex selection. The detailed structure of the *i*th layer in the Transformer decoder is depicted in the right part, where F denotes the output of the Transformer encoder, C_i , C_{i+1} represent content queries from different layers, while Q_i , Q_{i+1} denote room queries from different layers.

3.1 Floorplan Representation

In general, since the fundamental structure of a floorplan comprises flat walls connected by angular corners, it is natural and intuitive to represent it with a sequence of 2D polygon vertices [19, 39]. As depicted in Figure 3(a), previous study [39] has adopted a corner-based sparse floorplan representation, where each polygon possesses a varying number of corners. However, the sparse supervision may lead to erroneous contour deformations and collapsed topology in room polygons where corners are either missing or biased. For better supervision of the polygon contours during training, we opt for a uniform sampling floorplan representation. Here, each polygon is represented by a fixed number of vertices rather than solely relying on corners. This strategy enables dense supervision along all room walls holistically, thereby bolstering the model's robustness against missing or biased turning corner vertices.

Considering a scene *S* comprising M_{gt} rooms, the floorplan could be represented as: $S = \{R_i\}_{i=1}^{M_{gt}}$, where R_i denotes the *i*th room. For a single room, we represent it with *N* vertices, which start from the upper-left corner and are arranged in clockwise order: $R_i = (\mathbf{v}_1^i, \mathbf{v}_2^i, ..., \mathbf{v}_j^i, ..., \mathbf{v}_N^i)$, where \mathbf{v}_j^i denotes the *j*th vertex. Each vertex is represented by both its position and corner label: $\mathbf{v}_j^i = (\mathbf{p}_j^i, l_j^i)$, where $\mathbf{p}_j^i = (x_j^i, y_j^i)$ denotes position coordinate and $l_j^i \in \{0, 1\}$ indicates whether it is a corner (1) or not (0). To obtain the vertex set, we uniformly sample *N* vertices along the room contour at equal intervals and set all vertex labels to 0. Subsequently, for each corner, we select the sampling vertex closest to the corner and replace it with the actual corner, marking it with the label 1. With all corners included, the floorplan is represented by the vertex set.



Fig. 3: (a) Illustration of floorplan representation including the sparse corner-based (left) and our dense uniform sampling representation (right). Valid contour vertices (outlined) and corner vertices (filled) with supervision during training are colored according to the room on the left. (b) Illustration of the self-attention variants including the room-aware self-attention and vanilla self-attention. Our room-aware self-attention is a combination of intra-room and inter-room self-attention, which works among different vertices in a single room and among different rooms. And the vanilla self-attention performs on the flattened queries.

In this novel representation, dense supervision is achieved with a fixed number of vertices during training. Additionally, the introduction of room-aware query initialization and the utilization of angles are also enabled. Because in the corner-based sparse representation, room-aware query initialization becomes unattainable due to the unknown corner number, while angles become highly inaccurate if a certain corner is biased or missing.

3.2 Encoder

By projecting an indoor 3D point cloud along the gravity direction, PolyRoom takes the density map generated as input. It then utilizes a CNN backbone [17] to extract multi-scale image features. Then these features from different scales are flattened, concatenated, and added with positional encoding before serving as input to the Transformer encoder [36]. Within the encoder, each layer comprises a multi-scale deformable self-attention layer [43] and a feed-forward network. Finally, the features output by the encoder interact with queries in the cross-attention layer of the decoder.

3.3 Decoder

For the decoder structure, as depicted in Figure 2, each layer comprises a roomaware self-attention module which combines intra-room self-attention and interroom self-attention, along with a multi-scale deformable cross-attention module and a feed-forward network. Within each layer, queries initially interact within themselves in the room-aware self-attention module, followed by interaction with the encoder output in the cross-attention module. Ultimately, vertex coordinates are updated through the iterative refinement of room queries layer by layer. For the queries (q), they consist of content queries $C \in \mathbb{R}^{M \times N \times d}$ which are randomly initialized and positional queries P = MLP(PE(Q)). Here, M denotes the maximum polygon sequence length, d represents the dimension of decoder embedding, MLP refers to a simple multi-layer perception network and PE signifies the positional encoding, which calculates the d-dimensional sinusoidal positional embedding of the room queries $Q \in \mathbb{R}^{M \times N \times 2}$.

In vanilla self-attention, queries are flattened to interact with each other $(q \in \mathbb{R}^{MN \times d})$, which is constrained by memory consumption with increasing query length. Inspired by the self-attention variant in [28], we introduce the room-aware self-attention, where self-attention is performed along both the intraroom dimension $(q \in \mathbb{R}^{N \times Md})$ and the inter-room dimension $(q \in \mathbb{R}^{M \times Nd})$ (illustrated in Figure 3(b)). Based on the two-level floorplan representation, the room-aware self-attention focuses on relationships within individual rooms as well as among different rooms. It reduces the computational complexity from $O((MN)^2)$ to $O(M^2 + N^2)$, and results in less memory consumption. Moreover, it also achieves better reconstruction results than vanilla self-attention.

After the room-aware self-attention, the queries interact with the features output by the Transformer encoder in the cross-attention module. Following this interaction, vertex offsets $(\delta x, \delta y)$ are predicted from content queries with an MLP. Upon adding the vertex offsets, room queries are updated. With room queries explicitly represented as polygon vertex coordinates and updated after each decoder layer, room polygons are refined layer by layer. Finally, the polygon vertices are selected to extract the final vectorized floorplan.

3.4 Room-Aware Query Initialization

With the aforementioned advantages, the uniform sampling representation also presents a challenge due to longer vertex sequences, making it more difficult for randomly initialized room queries to be accurately learned. This can result in inaccuracies like self-interaction, overlapping, and incorrect sequence orders, as the queries lack a direct correlation to the inputs. Consequently, a room-aware query initialization module becomes imperative. Building upon the new floorplan representation, we propose a method to address the aforementioned challenge by leveraging instance segmentation to generate initial shapes for room queries.

In detail, the objective of room-aware query initialization is to provide a reasonable estimation of room contours for these queries. To achieve this, room instances are predicted and then represented as room queries. Specifically, the pre-trained instance segmentation network first predicts several room masks. With the room polygons, starting from the upper-left corner and moving clockwise, we uniformly sample N vertices along the polygon contour at equal intervals. In this way, room queries are explicitly represented as $Q \in \mathbb{R}^{M \times N \times 2}$ with vertex coordinates. If the number of room masks is less than M, we use random initial values to pad the room queries.

3.5 Floorplan Extraction

Room queries updated in the Transformer decoder correspond to the predicted polygon contours. For corner prediction, vertex labels are predicted with a feedforward network from the decoder output. In practice, vertices around corners tend to exhibit higher prediction probabilities to be corners, and multiple vertices may be identified as corners based solely on a probability threshold. To prevent redundant corners from being selected, we combine prediction probabilities and geometric properties for vertex selection to extract the final floorplan.

First, we select vertices predicted to be corners. Then, based on the prior knowledge of indoor buildings that room corner angles are typically not flat or small acute angles, we also include vertices based on an angle threshold. Taken together, the conditions for vertex selection can be expressed as:

$$pro_j > t_{pro}, \quad |ang_j| < t_{ang},$$
 (1)

where pro_j is the probability predicted to be corners, t_{pro} is the threshold for probability, ang_j indicates the cosine value of the vertex angle and t_{ang} is the threshold for the cosine value.

Given the structured and compact nature of floorplan polygons, we utilize a polygonization algorithm [12] to further select vertices. This ensures that the final vertex sets are minimal, exhibiting robust structural integrity without redundancy, while faithfully preserving the geometry of the original polygon. Subsequently, the final floorplan is extracted.

3.6 Loss Functions

Bipartite Matching. To make the network end-to-end trainable, it is supposed to find a bipartite matching between the predicted queries and the ground truth. With the ordered and dense representation, we conduct the bipartite matching totally based on polygon geometry. In detail, we compute distances between the sampling vertices of the predictions and the ground truth, and define the matching cost accordingly as follows:

$$C = \sum_{i=1}^{M_{gt}} \sum_{j=1}^{N} d(p_j^i, p_j^{\sigma(i)}),$$
(2)

where C is the matching cost, p_j^i and $p_j^{\sigma(i)}$ are vertices in the ground truth polygon and the corresponding matching polygon, d is the L_1 distance between two vertices. All possible arrangements of polygon matching are calculated, and an optimal bipartite matching is assigned with the minimum matching distance.

Total Loss. After matching, losses can be calculated between the predictions and ground truth. The overall loss function is defined as:

$$L = \sum_{i=1}^{M} (\lambda_{cls} L^{i}_{cls} + \lambda_{coord} L^{i}_{coord} + \lambda_{ras} L^{i}_{ras} + \lambda_{ang} L^{i}_{ang}),$$
(3)

where L_{cls} , L_{coord} , L_{ras} , and L_{ang} represent the cross-entropy classification loss, L1 coordinate loss, rasterization loss and angle loss respectively. λ_{cls} , λ_{coord} , λ_{ras} , and λ_{ang} are the corresponding loss weights. The rasterization loss is also used in RoomFormer [39], measuring differences between the reconstructed and GT shapes through differentiable rendering. The angle loss calculates angle differences to better supervise the densely sampled vertex positions:

$$L_{ang}^{i} = \frac{1}{N} \sum_{j=1}^{N} |ang_{j}^{i} - ang_{j}^{\sigma(i)}|, i \le M_{gt},$$
(4)

where ang_j^i and $ang_j^{\sigma(i)}$ represent the cosine value of the vertex angle in the ground truth and matching polygon respectively.

4 Experiments

4.1 Datasets and Evaluation Metrics

Datasets. We evaluate our proposed method in two popular indoor environment datasets, Structured3D [41] and SceneCAD [1]. Structured3D is a synthetic dataset with 3D structure annotations. There are 3500 houses in the dataset (3000/250/250 for training/validation/test) containing rich floorplan types. SceneCAD provides CAD-based representations of scanned 3D environments [10] from commodity RGB-D sensors, and floorplan annotations are extracted from them. The data of both the two datasets are projected along the vertical axis into 256×256 top-view point-density images as inputs.

Evaluation Metrics. For evaluation metrics, we adopt the same criteria as previous works [8, 39], including three-level metrics: Room, Corner, and Angle. Each level metric comprises Precision, Recall, and F1 score. Additionally, for the Room level metric, Intersection over Union (IoU) is reported.

4.2 Implementation Details

We utilize ResNet-50 [17] as our backbone and pre-train Mask2Former [9] on each dataset for 200 epochs as our instance segmentation network. We set the maximum number of rooms as 20 and the number of sampling vertices as 40. For the loss weights, we assign the values of λ_{cls} , λ_{coord} , λ_{angle} , and λ_{ras} as 2, 5, 1, and 1. For vertex selection, t_{pro} is set to 0.01, t_{ang} is set to $\sqrt{3}/2$, and the DP threshold is set to 4 for the Structured3D dataset and 10 for the SceneCAD dataset. During training, the model is trained on Structured3D for 500 epochs with a learning rate of 2e-4 and on SceneCAD for 400 epochs with a learning rate of 5e-5. All experiments are optimized by Adam [23] and implemented in pytorch with a batch size of 40 on an NVIDIA RTX 3090 GPU.

Table 1: Quantitative evaluation results on Structured3D [41] dataset. Results of previous works are taken from [39] and [34]. The colors orange and cyan mark the top-two results, where * indicates the use of our vertex selection approach.

Method			Room		Corner			Angle		
	Presented at	Prec.	Rec.	F1	Prec.	Rec.	F1	Prec.	Rec.	F1
DP	-	96.0	95.0	95.5	82.8	78.6	80.7	57.3	54.3	55.8
Floor-SP [7]	ICCV19	89.	88.	88.	81.	73.	76.	80.	72.	75.
MonteFloor [33]	ICCV21	95.6	94.4	95.0	88.5	77.2	82.5	86.3	75.4	80.5
HAWP [38]	CVPR20	77.7	87.6	82.3	65.8	77.0	70.9	59.9	69.7	64.4
LETR [37]	CVPR21	94.5	90.0	92.2	79.7	78.2	78.9	72.5	71.3	71.9
HEAT [8]	CVPR22	96.9	94.0	95.4	81.7	83.2	82.5	77.6	79.0	78.3
RoomFormer [39]	CVPR23	97.9	96.7	97.3	89.1	85.3	87.2	83.0	79.5	81.2
RoomFormer [*] [39]	CVPR23	97.8	96.2	97.0	91.6	83.8	87.5	86.1	78.9	82.3
SLIBO-Net [34]	NeurIPS23	99.1	97.8	98.4	88.9	82.1	85.4	87.8	81.2	84.4
PolyRoom (Ours)	-	98.9	97.7	98.3	94.6	86.1	90.2	89.3	81.4	85.2

Table 2: Quantitative evaluation results on SceneCAD [1] dataset. The colors orange and cyan mark the top-two results, where * indicates the use of our vertex selection approach.

	Room		Corner		Angle		
Method	IoU	Prec.	Rec.	F1	Prec.	Rec.	F1
DP	90.1	62.4	84.3	71.7	24.2	32.3	27.7
Floor-SP	91.6	89.4	85.8	87.6	74.3	71.9	73.1
HEAT	84.9	87.8	79.1	83.2	73.2	67.8	70.4
RoomFormer	91.7	92.5	85.3	88.8	78.0	73.7	75.8
RoomFormer [*]	91.3	94.4	84.7	89.2	79.6	73.4	76.4
PolyRoom (Ours)	92.8	96.8	86.1	91.2	81.7	74.5	78.0

4.3 Comparison with State of the Arts

Quantitative Evaluations. We compare our method with eight other approaches, including DP [12], Floor-SP [7], MonteFloor [33], HAWP [38], LETR [37], HEAT [8], RoomFormer [39], and SLIBO-Net [34]. DP simplifies room segments with DP [12] polygonization algorithm. Floor-SP and MonteFloor are optimization algorithms with room segmentation. HAWP and LETR are general methods for wireframe parsing and line segment detection. HEAT, RoomFormer, and SLIBO-Net are Transformer-based methods for floorplan reconstruction.

On Structured3D (Table 1), PolyRoom is compared with all the comparative methods, and on SceneCAD (Table 2), it is compared with four representative ones, including DP, Floor-SP, HEAT, and RoomFormer. On both datasets, Poly-Room outperforms all other comparative methods. Notably, on Structured3D, PolyRoom enhances the F1 score of corner and angle by 4.8% and 0.8% respectively compared to the SLIBO-Net which relies on the Manhattan assumption, while maintaining a similar result in the F1 score of the room. In comparison to RoomFormer, PolyRoom surpasses it across all three metrics significantly. On SceneCAD, PolyRoom achieves notable improvements, with increases of 1.1%, 2.0%, and 1.6% for Room IoU, Corner F1 score, and Angle F1 score respectively.



Fig. 4: Qualitative evaluations on Structured3D [41]. PolyRoom performs with fewer missing rooms, more correct room sequences, and improved room details.



Fig. 5: Qualitative evaluations on SceneCAD [1].

Qualitative Evaluations. Some results on Structured3D and SceneCAD are shown in Figure 4 and Figure 5 for qualitative comparison, it can be observed that MonteFloor and Floor-SP exhibit incorrect shapes of rooms, HEAT encounters issues with missing corners and edges, resulting in unclosed room polygons. RoomFormer suffers from missing rooms and inaccurate polygon sequences, such as self-intersection and overlap of polygons. In contrast, PolyRoom delivers more complete reconstruction results with fewer missing rooms, more accurate polygon sequences, and improved reconstruction details.

Cross-data Generalization. We evaluate the cross-data generalization ability of PolyRoom across different datasets. Our model is trained on the Structured3D training set and then tested on the SceneCAD validation set. We compare the results with two end-to-end competing floorplan reconstruction methods: HEAT and RoomFormer. The findings presented in Table 3 demonstrate that Poly-Room exhibits superior generalization compared to other methods, owing to the proposed room-aware query initialization and dense supervision mechanisms.

4.4 Ablation Studies

Different Components. Table 4 illustrates the impacts of different components on the final results. Without uniform sampling representation, room-aware query initialization introduces room initial shapes with varying vertex numbers

Table 3: Cross-data generalization. The best results are in **bold** font, the same as in the following tables.

	Room Corner				Angle				
Method	IoU	Prec.	Rec.	F1	Prec.	Rec.	F1		
HEAT	52.5	50.9	51.1	51.0	42.2	42.0	41.6		
RoomFormer	74.0	56.2	65.0	60.3	44.2	48.4	46.2		
PolyRoom (Ours)	85.2	77.8	79.9	78.9	59.4	61.7	60.6		

Table 4: Ablation study results of different components of PolyRoom on Structure3D [41].

	Ro	om	Cor	Corner		gle
Settings	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.
Baseline	98.9	97.7	94.6	86.1	89.3	81.4
w/o Uniform sampling representation	86.7	70.4	78.2	61.2	71.8	56.2
w/o Room-aware query initialization	96.9	96.4	89.3	84.6	82.7	78.4
w/o Room-aware self-attention	98.8	97.4	94.1	85.8	89.0	81.2
w/o Vertex selection	98.7	97.6	81.4	87.9	72.6	77.9

Table 5: Analysis on layer number and vertex number of PolyRoom on Structure3D [41].

Se	Settings		Room		mer	Angle	
Layer	Vertex	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.
6	40	98.9	97.7	94.6	86.1	89.3	81.4
4	40	98.2	96.9	93.5	85.0	87.9	80.1
8	40	98.7	97.6	94.7	86.0	89.2	81.3
6	30	98.2	97.2	93.9	85.9	87.8	80.4
6	50	98.8	97.8	94.4	86.3	87.6	80.2

compared to the ground truth, posing a challenge in predicting accurate corner counts. Conversely, without room-aware query initialization, employing a single uniform sampling representation results in a higher number of vertices, hindering the accurate regression of vertex coordinates and leading to inferior performance. Additionally, room-aware self-attention exhibits superior performance compared to vanilla self-attention. Regarding vertex selection, it removes redundant vertices, yielding more accurate and structured results.

Vertex and Layer Number. In this experiment, we investigate the effects of vertex number and layer number on the final results. As depicted in Table 5, within a certain range, as the layer number and vertex number increase, the model performance increases, while a high layer number yields similar results, and a high vertex number results in a decrease in performance. To strike a balance between accuracy and training time, we select a layer number of 6 and a vertex number of 40 as our parameters.

Table 6: Robustness to initialization on Structure3D [41].

				Room		Corner		Angle	
Network	Backbone	Para.(M)	mAP	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.
Mask-RCNN	Res50	44.1	39.3	98.6	97.3	92.6	86.0	87.0	81.1
Mask2Former	Res50	44.0	41.6	98.6	97.4	93.0	85.9	87.5	81.1
Mask2Former	Res101	63.0	44.3	98.7	97.4	94.0	86.0	88.3	81.2
Mask2Former Mask2Former	Swin-S Swin-L	68.7 215.5	$45.3 \\ 48.0$	98.7 98.9	97.4 97.7	94.2 94.6	86.0 86.1	88.8 89.3	81.2 81.4

Table 7: Ablation on the self-attention variants on Structure3D [41]. The GPU memory is tested with a batch size of 20, as the model with vanilla self-attention runs out of memory with a batch size of 40.

Self-attention		Room		Cor	mer	Angle	
	Memory	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.
Vanilla	18707 M	98.8	97.4	94.1	85.8	89.0	81.2
Single intra-room	11071M	98.2	96.9	93.3	84.9	88.2	80.4
Single inter-room	10947M	98.6	97.3	94.3	85.7	88.6	80.8
Room-aware	$11957 \mathrm{M}$	98.9	97.7	94.6	86.1	89.3	81.4

Table 8: Loss function and vertex selection analysis on Structure3D [41].

	Ro	om	Cor	ner	Angle	
Settings	Prec.	Rec.	Prec.	Rec.	Prec.	Rec.
Ours	98.9	97.7	94.6	86.1	89.3	81.4
w/o Angle loss	98.1	96.5	92.6	85.2	85.0	78.3
w/o Angle threshold w/o Polygonization	98.7 98.8	97.6 97.7	93.8 92.6	86.3 87.0	88.2 86.7	81.1 81.5

Robustness to Initialization. To further investigate the impact of initialization on the final results, we compare different instance segmentation networks, including the simple Mask-RCNN [16] and state-of-the-art Mask2Former [9] with different backbones. Table 6 presents the reconstruction results based on different segmentation qualities (measured by mean Average Precision (mAP)) and model sizes (measured by parameters (Para.)). Enhanced instance segmentation leads to improved room-aware query initialization, resulting in superior reconstruction results. Additionally, the final results also demonstrate robustness to segmentation results of different qualities.

Self-attention Variants. In Table 7, we conduct a detailed ablation experiment to validate the effectiveness of room-aware self-attention. It demonstrates much less training memory compared to vanilla self-attention while also achieving superior performance. Furthermore, compared to using intra-room and interroom self-attention separately, room-aware self-attention proves to be more suitable for our floorplan representation and achieves better performance.



Fig. 6: Failure cases. (Left) Our model misses rooms because of poor initialization. (Right) Our model misses details when dealing with complex structures. (Blue and orange boxes mark areas where reconstruction fails.)

Loss Functions and Vertex Selection. In this experiment, we evaluate the necessity of angle loss, angle threshold, and polygonization in vertex selection. Angle loss serves as an auxiliary loss focusing on angle perspective to better supervise vertex positions. Table 8 illustrates that the introduction of angle loss improves reconstruction results, particularly benefiting angle prediction. Concerning vertex selection, both angle threshold and polygonization demonstrate a positive impact in reducing redundant vertices.

4.5 Analysis of Failure Cases

Despite exhibiting improved robustness with segmentation networks of varying qualities, our model still encounters difficulty in accurately predicting the complete floorplan when certain room instances are missing from the segmentations (left in Figure 6). Moreover, when dealing with complex indoor structures, PolyRoom could overlook details such as thin walls, which can be attributed to limitations in segmentation results and raw data quality (right in Figure 6).

5 Conclusion

In this paper, we present PolyRoom, a room-aware Transformer for floorplan reconstruction. Utilizing uniform sampling representation, room-aware query initialization, and room-aware self-attention, PolyRoom offers significant advancements in the field. With room-aware query initialization, we establish a direct link between room semantics and room query positions, enhancing the reconstruction accuracy. Subsequently, room queries are progressively refined in the decoder layer by layer, benefiting from dense supervision throughout. The final floorplan is derived through vertex selection, incorporating both probabilistic considerations and geometric constraints. Moreover, the adoption of room-aware self-attention not only reduces memory consumption but also improves reconstruction quality. Both qualitative and quantitative evaluations demonstrate the superiority of PolyRoom over other existing methods while with better generalization. Our comprehensive ablation study further validates the efficacy of our proposed design choices. We hope our work will inspire further research on reconstruction methods with polygon sequences.

Acknowledgements

We thank the authors of RoomFormer, HEAT and MonteFloor for providing experimental results, especially Yuanwen Yue for his kind help with the experiments. This work was supported by the National Natural Science Foundation of China (No. U22B2055, 62273345, and 62373349), the Beijing Natural Science Foundation (No. L223003), and the Key R&D Project in Henan Province (No. 231111210300).

References

- Avetisyan, A., Khanova, T., Choy, C., Dash, D., Dai, A., Nießner, M.: SceneCAD: Predicting object alignments and layouts in RGB-D scans. In: European Conference on Computer Vision (ECCV). pp. 596–612 (2020)
- Browne, C.B., Powley, E., Whitehouse, D., Lucas, S.M., Cowling, P.I., Rohlfshagen, P., Tavener, S., Perez, D., Samothrakis, S., Colton, S.: A survey of monte carlo tree search methods. IEEE Transactions on Computational Intelligence and AI in Games 4(1), 1–43 (2012)
- Cabral, R., Furukawa, Y.: Piecewise planar and compact floorplan reconstruction from images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 628–635 (2014)
- Castrejon, L., Kundu, K., Urtasun, R., Fidler, S.: Annotating object instances with a polygon-RNN. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4485–4493 (2017)
- Chen, C., Wang, R., Vogel, C., Pollefeys, M.: F3Loc: Fusion and filtering for floorplan localization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18029–18038 (2024)
- Chen, J., Deng, R., Furukawa, Y.: PolyDiffuse: Polygonal shape reconstruction via guided set diffusion models. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 1863–1888 (2023)
- Chen, J., Liu, C., Wu, J., Furukawa, Y.: Floor-SP: Inverse CAD for floorplans by sequential room-wise shortest path. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 2661–2670 (2019)
- Chen, J., Qian, Y., Furukawa, Y.: HEAT: Holistic edge attention Transformer for structured reconstruction. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3856–3865 (2022)
- Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask Transformer for universal image segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1280–1289 (2022)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scan-Net: Richly-annotated 3D reconstructions of indoor scenes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2432–2443 (2017)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16×16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (ICLR) (2020)
- Douglas, D.H., Peucker, T.K.: Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. Cartographica 10(2), 112– 122 (1973)

- 16 Y. Liu et al.
- Fan, Z., Zhu, L., Li, H., Chen, X., Zhu, S., Tan, P.: FloorPlanCAD: A large-scale CAD drawing dataset for panoptic symbol spotting. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10108–10117 (2021)
- Favreau, J.D., Lafarge, F., Bousseau, A., Auvolat, A.: Extracting geometric structures in images with Delaunay point processes. IEEE Transactions on Pattern Analysis and Machine Intelligence 42(4), 837–850 (2019)
- Han, J., Liu, Y., Rong, M., Zheng, X., Shen, S.: FloorUSG: Indoor floorplan reconstruction by unifying 2D semantics and 3D geometry. ISPRS Journal of Photogrammetry and Remote Sensing 196, 490–501 (2023)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: IEEE International Conference on Computer Vision (ICCV). pp. 2980–2988 (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural Computation 9(8), 1735–1780 (1997)
- Hu, Y., Wang, Z., Huang, Z., Liu, Y.: PolyBuilding: Polygon Transformer for building extraction. ISPRS Journal of Photogrammetry and Remote Sensing 199, 15–27 (2023)
- Ibrahem, H., Salem, A., Kang, H.S.: ST-RoomNet: Learning room layout estimation from single image through unsupervised spatial transformations. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3376–3384 (2023)
- Ikehata, S., Yang, H., Furukawa, Y.: Structured indoor modeling. In: IEEE International Conference on Computer Vision (ICCV). pp. 1323–1331 (2015)
- Jiang, Z., Xiang, Z., Xu, J., Zhao, M.: LGT-Net: Indoor panoramic room layout estimation with geometry-aware Transformer network. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1644–1653 (2022)
- 23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR) (2015)
- Lazarow, J., Xu, W., Tu, Z.: Instance segmentation with mask-supervised polygonal boundary Transformers. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4372–4381 (2022)
- Lee, C.Y., Badrinarayanan, V., Malisiewicz, T., Rabinovich, A.: RoomNet: Endto-end room layout estimation. In: IEEE International Conference on Computer Vision (ICCV). pp. 4875–4884 (2017)
- Li, J., Chan, C.L., Le Chan, J., Li, Z., Wan, K.W., Yau, W.Y.: Cognitive navigation for indoor environment using floorplan. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 9030–9037. IEEE (2021)
- Liang, J., Homayounfar, N., Ma, W.C., Xiong, Y., Hu, R., Urtasun, R.: Poly-Transform: Deep polygon transformer for instance segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9128–9137 (2020)
- Liao, B., Chen, S., Zhang, Y., Jiang, B., Zhang, Q., Liu, W., Huang, C., Wang, X.: MapTRv2: An end-to-end framework for online vectorized HD map construction. arXiv preprint arXiv:2308.05736 (2023)
- Liu, H., Cao, C., Ye, H., Cui, H., Gao, W., Wang, X., Shen, S.: Lightweight structured line map based visual localization. IEEE Robotics and Automation Letters 9(6), 5182–5189 (2024)

- Liu, J., Ding, H., Cai, Z., Zhang, Y., Satzoda, R.K., Mahadevan, V., Manmatha, R.: PolyFormer: Referring image segmentation as sequential polygon generation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18653–18663 (2023)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin Transformer: Hierarchical vision Transformer using shifted windows. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9992–10002 (2021)
- Peng, S., Jiang, W., Pi, H., Li, X., Bao, H., Zhou, X.: Deep snake for real-time instance segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8530–8539 (2020)
- 33. Stekovic, S., Rad, M., Fraundorfer, F., Lepetit, V.: MonteFloor: Extending MCTs for reconstructing accurate large-scale floor plans. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16014–16023 (2021)
- Su, J.W., Tung, K.Y., Peng, C.H., Wonka, P., Chu, H.K.: SLIBO-Net: Floorplan reconstruction via slicing box representation with local geometry regularization. In: Advances in Neural Information Processing Systems (NeurIPS). pp. 48781–48792 (2023)
- Sun, C., Hsiao, C.W., Sun, M., Chen, H.T.: HorizonNet: Learning room layout with 1D representation and pano stretch data augmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1047–1056 (2019)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems (NIPS) (2017)
- 37. Xu, Y., Xu, W., Cheung, D., Tu, Z.: Line segment detection using Transformers without edges. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4255–4264 (2021)
- Xue, N., Wu, T., Bai, S., Wang, F., Xia, G.S., Zhang, L., Torr, P.H.: Holisticallyattracted wireframe parsing. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2785–2794 (2020)
- 39. Yue, Y., Kontogianni, T., Schindler, K., Engelmann, F.: Connecting the dots: Floorplan reconstruction using two-level queries. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 845–854 (2023)
- Zhang, F., Xu, X., Nauata, N., Furukawa, Y.: Structured outdoor architecture reconstruction by exploration and classification. In: IEEE/CVF International Conference on Computer Vision (ICCV). pp. 12407–12415 (2021)
- Zheng, J., Zhang, J., Li, J., Tang, R., Gao, S., Zhou, Z.: Structured3D: A large photo-realistic dataset for structured 3D modeling. In: European Conference on Computer Vision (ECCV). pp. 519–535 (2020)
- Zheng, Z., Li, J., Zhu, L., Li, H., Petzold, F., Tan, P.: GAT-CADNet: Graph attention network for panoptic symbol spotting in CAD drawings. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11737– 11746 (2022)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable Transformers for end-to-end object detection. In: International Conference on Learning Representations (ICLR) (2020)