

Polygon Detection for Room Layout Estimation using Heterogeneous Graphs and Wireframes

David Gillsjö Gabrielle Flood Kalle Åström
Centre for Mathematical Sciences, Lund University, Sweden
{david.gillsjo,gabrielle.flood,kalle.astrom}@math.lth.se

Abstract

This paper presents a neural network based semantic plane detection method utilizing polygon representations. The method can for example be used to solve room layout estimations tasks and is built on, combines and further develops several different modules from previous research. The network takes an RGB image and estimates a wireframe as well as a feature space using an hourglass backbone. From these, line and junction features are sampled. The lines and junctions are then represented as an undirected graph, from which polygon representations of the sought planes are obtained. Two different methods for this last step are investigated, where the most promising method is built on a heterogeneous graph transformer. The final output is in all cases a projection of the semantic planes in 2D. The methods are evaluated on the Structured3D dataset and we investigate the performance both using sampled and estimated wireframes. The experiments show the potential of the graph-based method by outperforming state of the art methods in Room Layout estimation in the 2D metrics using synthetic wireframe detections.

1. Introduction

With all the advanced methods in Computer Vision today, it is possible to perform Structure from Motion and 3D reconstruction in many ways. Depending on your application a generic point-based method might work well, but in other cases it may be beneficial to model projections as other geometric objects. In this paper we present

Acknowledgments: This work was supported by the Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation (KAW), the strategic research project ELLIIT, the ADACORSA project with funding from ECSEL JU in the H2020 Framework Programme (H2020/2014-2020) and National Authorities (GA 876019). The compute was enabled by the supercomputing resource Berzelius provided by National Supercomputer Centre at Linköping University and KAW.



Figure 1: The model takes an image and detects a wireframe which it uses to produce a set of polygons corresponding to a Room Layout.

a method for detecting the projections of planes in an image, where each projection may be represented as a polygon. The input to the suggested system is an RGB image and the first step in the pipeline is to detect a wireframe consisting of line segments in this image, as shown in Figure 1. We try two different methods for extracting polygons from the estimated wireframe. First, we try a cycle generating method that serves as a naive baseline which heuristically finds all possible polygons from the wireframe and classifies the polygons using a Convolutional Neural Network (CNN). Secondly, we develop our proposed method based on a Heterogeneous Graph Transformer (HGT) network to utilize the wireframe structure for polygon classification. Both methods share architecture for the backbone and final classifier CNN. An overview of the developed systems is shown in Figure 2a.

The experiments in this paper focus on the Room Layout Estimation task, where each wall, floor and ceiling is represented as a plane polygon in 3D. Even though our system does not produce any depth information, we can still compare the 2D projections of the room layout. We show that given good enough line segment detections, the model successfully predicts a Room Layout projection with per-

formance matching the state of the art models. While there certainly are other tasks that fit the assumptions made with this model, there are good datasets for Room Layout Estimation to use when training the model. This means that our plane polygon detector will be evaluated on this type of tasks, but the usage is not limited to this and may be extended in the future.

Previous work in Room Layout Estimation that use Neural Network (NN) models produce an intermediate representation that is later on used in an heuristic optimization method to generate the Room Layout. For example Lin *et al.* [11] predict a semantic segmentation which is later used to optimize intersection lines heuristically. Other methods [18, 32, 26, 8] predict edges and then optimize. The most recent work of [27] use a combination of plane, depth and vertical line detections to heuristically fuse and generate a Room Layout. Most of these methods also place strong priors on the room shape, for example requiring cuboid shape or fixed angles between planes. We present a method in which the NN model reasons jointly over line predictions to directly estimate a Room Layout with the only constraint being that the room is a set of intersecting planes. To the best of our knowledge, this is the first end-to-end NN model to directly output a Room Layout without cuboid requirements.

The contributions of this paper are:

- A novel Heterogeneous Graph Neural Network (HGNN) model¹ and a polygon sampling model for estimating Room Layouts from line detections. The model does not require any post-processing heuristics.
- Results on the dataset Structured3D using generated and learned line proposals. This study gives insights about the importance of the graph structure of the problem.
- A quantitative comparison with state of the art models in Room Layout estimation.

2. Related Work

Plane Instance Detection is a related task where the goal is to detect each plane in the image and provide a pixel mask and 3D parameter estimate for each plane instance with the camera as origin. The planar objects are foreground objects so this task does not deal with occlusions, contrary to Room Layout estimation. The state of the art models are all NN models [30, 28, 13, 14, 21] with impressive performance on instance level.

Wireframe Parsing. As input for our model we use line segment detections from wireframe parsing [22, 23, 24, 25]. These methods predict connected line segments without any

semantic meaning. The most common pipeline is to start from junction proposals and then work toward line segments, but some works like [16] directly use line predictions together with a Graph Neural Network. Our previous work [3] predicts a semantic wireframe based on room geometry using a Graph Convolutional Network, but this does not account for plane instances.

Room Layout Estimation has been studied in many forms. Hedau *et al.* [5] early on used the Manhattan World assumption [1] as prior for the room shape. Prior to NN models [19, 2, 31] used handcrafted features as a base, then did vanishing point detection and hypothesis generation. Recent approaches use CNNs, *e.g.* Mallya and Lazebnik [15] utilize structured edge detection forests with a CNN to predict an edge probability mask. Lin *et al.* [11] proposed an end-to-end CNN for pixelwise segmentation of the room image with post processing. DeepRoom3D [12] use an end-to-end CNN to predict a cuboid and RoomNet [10] directly predicts ordered keypoints in a room layout. A group of recent methods [18, 32, 26, 8] use CNNs to predict edges and then optimize for the Room Layout using geometric priors.

All these approaches have strong priors on the Room geometry. There have recently been some methods considering a general model without these constraints. Stekovic *et al.* [20] solve a discrete optimization problem over 3D polygons using both RGB and Depth information. Howard-Jenkins *et al.* [6] use plane detection to form a 3D model over a video sequence. Yang *et al.* [27] use a combination of plane, depth and vertical line detections to estimate a general Room Layout. The datasets mostly used are LSUN [29] and Hedau [5]. In this paper we, similarly to [27], will use Structured3D [33]. There is also much work done in Room Layout estimation for panorama images, but as this work focuses on perspective images we deem it to be outside of the scope.

3. Model

In this paper we evaluate two different models for finding semantic planes in images, which both share the same backbone. The models will be described in detail later but here, and in Figure 2a, we give an overview of the two different models. The Cycle Sampling Polygon-based Classifier (CSP) generates possible polygons from detected line segments and samples features from the polygon for classification. The Heterogeneous Graph-based Classifier (HGC) predicts plane centroids to use as anchors. These anchors are together with line and junction features encoded into a heterogeneous graph which is fed through the network. In the end, each plane node is classified using an MLP classifier.

¹<https://github.com/DavidGillsjo/polygon-HGT>.

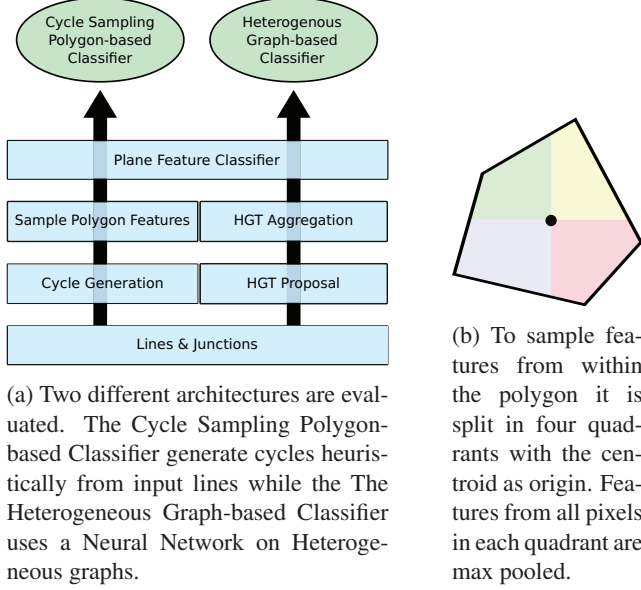


Figure 2: Model comparison and polygon sampling.

3.1. Backbone

The backbone is an Hourglass Neural Network [17] which takes a resized RGB image of size $3 \times 512 \times 512$ pixels as input and outputs a feature space \mathcal{F} with dimensions $128 \times 128 \times 256$, where 128 is spatial size and 256 is the number of channels. This feature space is used for sampling both junction, line and plane features.

3.2. Line and Junction Feature Sampling

The method for line and junction feature sampling is based on HAWP [25], but we add labels to the junctions to predict whether the junctions are corresponding to actual junctions in 3D or end a line segment due to occlusions. Similarly to SRW-net [3], if a junction has a 3D correspondence we call it *proper*, otherwise *false*.

Let $J_i = \{\mathbf{z}_i, \mathbf{f}_i\}$ denote a junction with index i and $\mathbf{J} = \{J_1, J_2, \dots, J_R\}$ a set of junctions. Each junction consists of the coordinates $\mathbf{z}_i \in \mathbb{R}_{[0,128]}^2$ and a feature vector $\mathbf{f}_i \in \mathbb{R}^{256}$. This feature vector is generated by

$$\mathbf{f}_i = \mathcal{J}(\mathcal{F}(\lfloor \mathbf{z}_i \rfloor)), \quad (1)$$

where \mathcal{J} is $\text{Conv2D}(256, 3) \rightarrow \text{ReLU}$. The notation for the convolution layer is $\text{Conv2d}(\text{output channels}, \text{kernel size})$.

Similarly, let $L_j = \{\mathbf{u}_j, \mathbf{v}_j, \mathbf{g}_j\}$ denote a line segment with index j consisting of end point coordinates $(\mathbf{u}_j, \mathbf{v}_j) \in \mathbb{R}_{[0,128]}^2$ and a feature vector $\mathbf{g}_j \in \mathbb{R}^{512}$. The endpoints are also junctions, and the feature vector is generated by a network head \mathcal{L}_{NN} and a sampling function \mathcal{L}_s where \mathcal{L}_{NN} is $\text{Conv2D}(128, 3) \rightarrow \text{ReLU}$. \mathcal{L}_s takes 16 sample points uniformly spread between endpoints $(\mathbf{u}_j, \mathbf{v}_j)$ and interpolates a feature vector of size 16×128 , this is down

sampled by max pooling to 4×128 and flattened so that $\mathcal{L}_s(L_j) \in \mathbb{R}^{512}$. In summary the vector \mathbf{g}_j is formed by taking

$$\mathbf{g}_j = \mathcal{L}_s(\mathcal{L}_{NN}(\mathcal{F}(\mathbf{u}_j, \mathbf{v}_j))). \quad (2)$$

3.3. Cycle Sampling Polygon-based Classifier

The first architecture (CSP) generates possible polygons from detected line segments and samples features from the polygon for classification. Given a set of connected line segments $\mathbf{L} = \{L_1, L_2, \dots, L_M\}$ we form an undirected graph \mathcal{G} with endpoints as vertices and form an edge for every detected line segment. From \mathcal{G} all possible polygons may be generated. This is done by generating cycles using the cycle basis of \mathcal{G} [9]. A cycle is a path of connected vertices $V = \{v_1, v_2, \dots, v_n, v_1\}$ where only the first and last vertex are equal. All simple cycles are not polygons, so we require each cycle to have at least three unique vertices and the generated polygon to have no intersecting boundaries. This is our baseline method and is further described in the supplementary material.

This gives a set of polygons $\mathbf{P} = \{P_1, P_2, \dots, P_K\}$. The Cycle Sampling Polygon-based Classifier proceeds to sample features from inside the polygon, similar to how object detectors like Mask-RCNN [4] samples from inside an object's bounding box. The sampling is done in three parts, first (i) all pixels of \mathcal{F} are processed by a CNN layer \mathcal{P}_{CNN} consisting of $\text{Conv2D}(128, 3) \rightarrow \text{ReLU}$. Then (ii) for each polygon P_k its centroid $\mathbf{c}_k \in \mathbb{R}_{[0,128]}^2$ is calculated. With \mathbf{c}_k as origin P_k is divided into four quadrants, see Figure 2b, where $\mathcal{P}_{CNN}(\mathcal{F})$ for each quadrant is max pooled. This yields a feature vector of size 4×128 which is flattened to be the polygon feature $\mathbf{h}_k \in \mathbb{R}^{512}$. The feature vector \mathbf{h}_k is then (iii) put through a 3 layer MLP for classification.

A major downside of the cycle sampling approach is that the number of possible cycles may grow exponentially with the number of edges and vertices. During inference it is also required to find all possible cycles to ensure that any valid polygon is found. However, at training time it is not necessary to find all cycles, so to speed up training only 20 polygons are sampled. We also add 10 positive examples and 10 negative examples of polygons from the annotations to stabilize training.

3.4. Heterogeneous Graph-based Classifier

The second architecture (HGC) is based on the Heterogeneous Graph Transformer [7]. An heterogeneous graph may consist of vertices and edges of different types. We naturally use the node types junction, line and plane, and also different edge types in the proposal network and the classifier network. Apart from the feature vectors for each entity we also encode the geometric position for the junctions, the midpoint for line segments, and the centroid for plane vertices. Figure 3 shows an illustration on how we

Polygon representation Heterogenous Graph representation




Figure 3: Illustration of how each polygon (left) may be represented with a heterogeneous graph (right) consisting of junction, line and plane vertices. In this illustration there are two types of edges, junction-line and line-plane.

represent a polygon as a heterogeneous graph.

An overview of the full architecture of the Heterogeneous Graph-based Classifier (HGC) can be seen in Figure 4. The junction and line features are sampled from \mathcal{F} according to Section 3.2, and geometric information is concatenated with the features. The proposal network takes the junctions \mathbf{J} and lines \mathbf{L} and generates proposal polygons \mathbf{P} . The polygons are together with the lines and junctions passed to the classifier network which labels each polygon.

3.4.1 Heterogeneous Graph Transformer

To jointly reason about planes, lines and junctions we use the Heterogeneous Graph Transformer (HGT) model [7] to perform message passing on the heterogeneous graph. Like a convolution operation in a CNN the message passing updates the current node features with a weighted sum of nearby node features. Instead of specifying a kernel we may choose a hop length. If the hop length is one, then all direct neighbours of node t are source nodes $N[t]$ for the message passing. A standard model to attention-based message passing in a homogeneous GNN for calculating layer b is

$$H^b[t] = \text{Aggregate}(\text{Message}(s) \cdot \text{Attention}(s, t)), \quad (3)$$

$\forall s \in N[t]$

where $\text{Message}(s)$ in the simplest case is $H^{b-1}[s]$, a neighbour node feature from previous layers. The Aggregate operation may be a sum, mean or other operation that suits the application.

For a heterogeneous graph each node t may have a different type $\tau(t)$ and each edge e a different edge type $\theta(e)$. The HGT model calculate different attention weights based on the meta relation $\langle \tau(t), \theta(e), \tau(s) \rangle$. For more details on the model, including aggregation and message construction, see the paper by Hu *et al.* [7].

3.4.2 Proposal Network

Let $\mathbf{P}_A = \{A_1, A_2, \dots, A_{25}\}$ denote a set of plane anchors placed uniformly on a 5×5 grid. Plane anchor A_k has coordinates $\mathbf{a}_k \in \mathbb{Z}_{[0,127]}^2$ and features \mathbf{d}_k generated by

$$\mathbf{d}_k = \mathcal{P}_p(\mathcal{P}_{CNN}(\mathcal{F}), \mathbf{a}_k), \quad (4)$$

where \mathcal{P}_{CNN} is the same as for CSP in Section 3.3 and \mathcal{P}_p is a multi-scale convolution operation with four $\text{Conv2D}(32, h, d)$ operations of kernel size $h = 1, 3, 3, 3$ and dilation $d = 0, 0, 1, 2$, respectively. The output of these four are concatenated to form the output of \mathcal{P}_p .

Now the graph \mathcal{G}^p for HGT is constructed. The features are

$$\bar{\mathbf{f}}_i = [\bar{\mathbf{z}}_i \quad \mathbf{f}_i], \quad \text{for } J_i \in \mathbf{J}, \quad (5)$$

$$\bar{\mathbf{g}}_j = [\frac{\bar{\mathbf{u}}_j + \bar{\mathbf{v}}_j}{2} \quad \mathbf{g}_j], \quad \text{for } L_j \in \mathbf{L}, \quad (6)$$

$$\bar{\mathbf{d}}_k = [\bar{\mathbf{a}}_k \quad \mathbf{d}_k], \quad \text{for } A_k \in \mathbf{P}_A, \quad (7)$$

where $\bar{\mathbf{z}} = \frac{\mathbf{z}}{128}$ to normalize the coordinates and likewise for $\bar{\mathbf{u}}, \bar{\mathbf{v}}, \bar{\mathbf{a}}$. As illustrated in Figure 5 the graph is passed through an HGT of l layers which yields a new graph \mathcal{G}_l^p with the same structure but updated features $\hat{\mathbf{f}}_i, \hat{\mathbf{g}}_j, \hat{\mathbf{d}}_k$. The graph \mathcal{G}_l^p is used to generate a plane proposal for each plane node. This problem is formulated as a link prediction task, where we for each plane anchor A_k predict edges to line nodes L_j . For each pair j, k a score is predicted as $\sigma(\hat{\mathbf{g}}_j^T W \hat{\mathbf{d}}_k)$, where $\sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function and W is a weight matrix learned during training. Then for each A_k a homogeneous graph with junctions as nodes and lines as edges is formed using every line j s.t. $\sigma(\hat{\mathbf{g}}_j^T W \hat{\mathbf{d}}_k) > \kappa$. The edge weight correspond to how likely it is that the line should be in the plane with anchor A_k . We find the cycle with Maximum average weight using an approximate greedy solution² and take it as plane proposal P_k .

To train the network a binary cross entropy loss is used with bounded biparte matching. First a biparte matching $\rho: \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ is found between all annotated line segments $\hat{L}_q \in \hat{\mathbf{L}}$ and detected line segment $L_j \in \mathbf{L}$ s.t.

$$\rho^* = \underset{\rho}{\operatorname{argmin}} \sum_q D(L_{\rho(q)}, \hat{L}_q),$$

$$D(L_j, \hat{L}_q) = \min(\|\mathbf{u}_j - \hat{\mathbf{u}}_q\|^2 + \|\mathbf{v}_j - \hat{\mathbf{v}}_q\|^2, \|\mathbf{u}_j - \hat{\mathbf{v}}_q\|^2 + \|\mathbf{v}_j - \hat{\mathbf{u}}_q\|^2)$$

if $D(L_{\rho(q)}, \hat{L}_q) < \alpha$, otherwise the match is excluded.

Let $\hat{\mathbf{P}}$ denote the set of G annotated ground truth planes with plane instances $\hat{P}_m \in \hat{\mathbf{P}}$ s.t. $\hat{P}_m = \{\hat{\mathbf{c}}_m, \hat{\mathbf{L}}_m\}$, where $\hat{\mathbf{L}}_m \subset \hat{\mathbf{L}}$ is the set of M_m line segments defining the polygon and $\hat{\mathbf{c}}_m$ the centroid of the polygon. Another biparte

²See supplementary material for details.

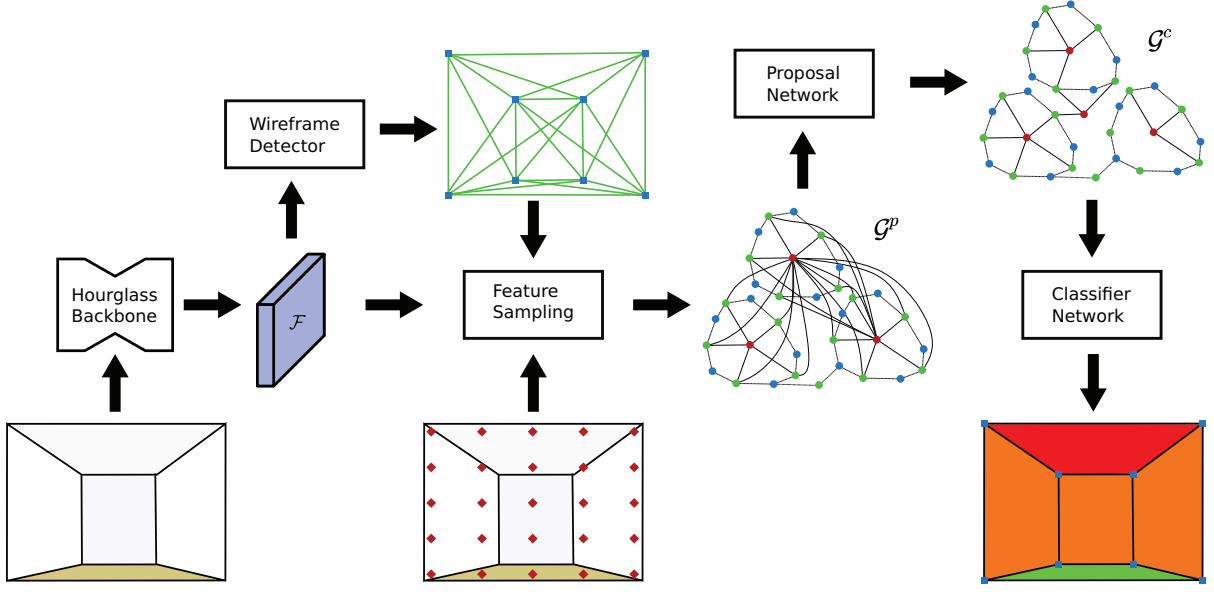


Figure 4: Simplified illustration of the Heterogeneous Graph-based Classifier architecture. The image is fed through a backbone to yield a feature space \mathcal{F} . The Wireframe input is then used together with plane centroid anchors to sample features from \mathcal{F} . These are passed to the proposal network and the classifier network as heterogeneous graphs.

mapping $\phi : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ is made between annotated planes and plane anchors $A_k \in \mathbf{A}$ s.t. the distance between centroids are minimized, i.e.

$$\phi^* = \underset{\phi}{\operatorname{argmin}} \sum_m \|\mathbf{a}_{\phi(m)} - \hat{\mathbf{c}}_m\|. \quad (8)$$

Now we may construct the loss function as

$$\mathcal{R}_{prop} = \frac{1}{G} \sum_{m=1}^G \frac{1}{M_m} \left[\sum_{L_j \in \Phi_m} \log \sigma \left(\hat{\mathbf{g}}_j^T W \hat{\mathbf{d}}_{\phi^*(m)} \right) + \sum_{L_j \notin \Phi_m} \log \left(1 - \sigma \left(\hat{\mathbf{g}}_j^T W \hat{\mathbf{d}}_{\phi^*(m)} \right) \right) \right],$$

where Φ_m are the detected lines matched with lines in the matched annotated plane \hat{P}_m , i.e. $\Phi_m = \{L_j \mid \rho^*(q) = j \text{ and } \hat{L}_q \in \hat{P}_m\}$.

3.4.3 Classifier Network

The classifier network takes the plane proposals $\mathbf{P} = \{P_1, P_2, \dots, P_K\}$ and constructs a graph \mathcal{G}^c by taking previous junction and line features from Equation (5) and (6) together with their structure from \mathcal{G}^p . In addition plane features \mathbf{d}_k^c are sampled from the plane proposal polygons by the same procedure as in the CSP model in Section 3.3. Just like in the proposal network, we add a normalized centroid

coordinate $\bar{\mathbf{c}}_k$ calculated from each polygon forming the final feature vector

$$\bar{\mathbf{d}}_k^c = [\bar{\mathbf{c}}_k \quad \mathbf{d}_k^c], \quad \text{for } P_k \in \mathbf{P}. \quad (9)$$

The graph \mathcal{G}^c is processed by a classifier HGT of l layers yielding a new graph \mathcal{G}_l^c with plane features $\hat{\mathbf{d}}_k^c$ for $P_k \in \mathbf{P}$. Final classification scores $\mathbf{s} \in \mathbb{R}_{[0,1]}^4$ for background, wall, floor and ceiling are given by network layers,

$$\mathbf{s}_k = \text{Softmax}(\text{MLP}(\hat{\mathbf{d}}_k^c)). \quad (10)$$

The loss function is a cross entropy loss from bounded bipartite matching. Using previous notation we find the mapping

$$\phi_c^* = \underset{\phi_c}{\operatorname{argmin}} \sum_m \|\mathbf{c}_{\phi_c(m)} - \hat{\mathbf{c}}_m\|, \quad (11)$$

if $\|\mathbf{c}_{\phi_c(m)} - \hat{\mathbf{c}}_m\| < \tau_c$. The loss function for annotated labels y_m is

$$\mathcal{R}_{cl} = \frac{1}{G} \sum_{m=1}^G -\mathbf{w}[y_m] \mathbf{s}_{\phi_c^*(m)}[y_m], \quad (12)$$

where $\mathbf{w}[y]$ is weight for class y and $\mathbf{s}_k[y]$ is what P_k scored for class y .

3.4.4 Joint Wireframe Estimation

This model is also trained to jointly output a Wireframe detection from the same backbone. The loss function \mathcal{R}_{wire}

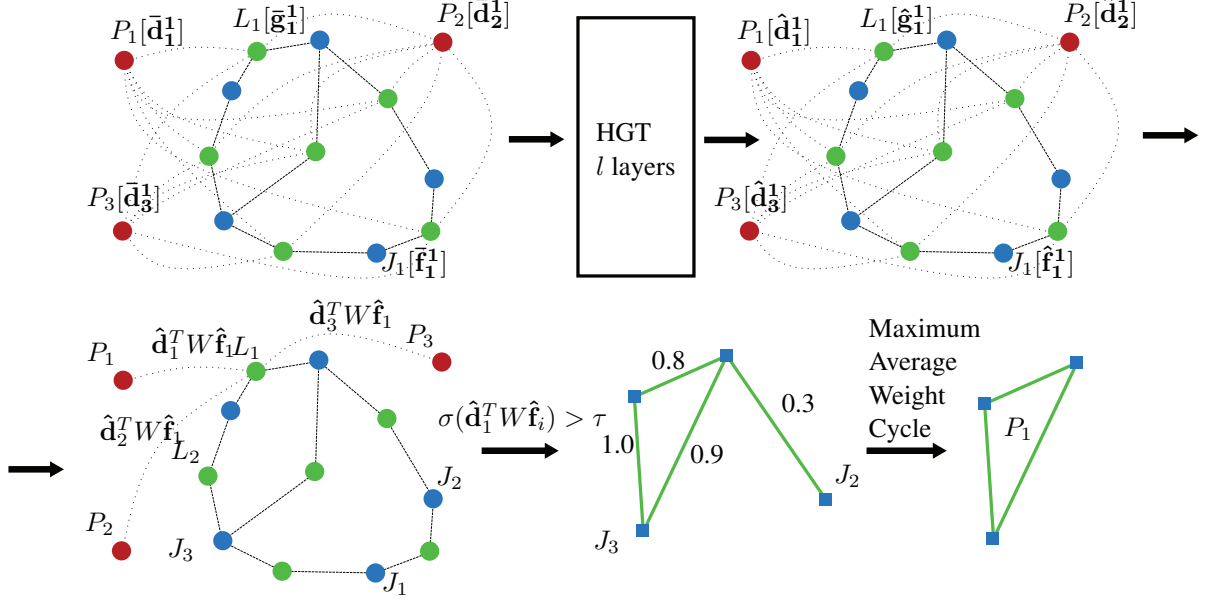


Figure 5: Illustration of the Proposal network in the Heterogeneous Graph-based Classifier. (i) All plane nodes are linked with all line nodes in the graph to form \mathcal{G}^p . (ii) The graph is fed to the HGT model and new features are computed, yielding \mathcal{G}_l^p . (iii) Link prediction between line and plane nodes is performed, edges with score $> \tau$ are transformed to a homogeneous graph and (iv) a polygon is found by finding the cycle with maximum average weight.

is defined by HAWP and so the total loss function for the network is

$$\mathcal{R} = \mathcal{R}_{wire} + \mathcal{R}_{prop} + \mathcal{R}_{cl}, \quad (13)$$

with weights for each term used during training.

3.4.5 Non-Maximum Suppression

Though the graph model seems to do just fine without a Non-Maximum Suppression (NMS) we implement one for the output plane polygons. A detected polygon $P_k \in \mathbf{P}$ is suppressed if there is another polygon $P_r \in \mathbf{P}$ with higher score $s_r > s_k$ and $\text{IoU}(P_r, P_k) > 0.05$.

4. Evaluation

To evaluate our model we cannot use all the Room Layout metrics since the models do not output depth information. We do use the Intersection over Union (IoU) error and Pixel error from Stekovic *et al.* [20]. For these metrics we find a one-to-one correspondence $\theta : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ mapping between the K predicted plane polygons \mathbf{P} and the M ground truth polygons $\hat{\mathbf{P}}$ s.t. $\theta(m) = k$ iff. $\text{IoU}(\hat{P}_m, P_k) \geq \text{IoU}(\hat{P}_m, P_i)$ for $i \neq k$ starting from the largest ground truth

polygon. Given this mapping we calculate

$$\epsilon_I[\text{IoU}] = \frac{2}{M+K} \sum_{m=1}^M \text{IoU}(\hat{P}_m, P_{\theta(m)}), \quad (14)$$

$$\epsilon_I[\text{PE}] = \frac{1}{|I|} \sum_{\mathbf{x} \in I} \text{PE}(\mathbf{x}), \quad (15)$$

where $\text{PE}(\mathbf{x}) = 0$ if the pixel is incorrectly matched, $\text{PE}(\mathbf{x}) = 1$ if correct matched and $|I|$ is the number of pixels. Furthermore, $\epsilon_I[\text{IoU}]$ and $\epsilon_I[\text{PE}]$ are calculated for each image I and then averaged to form $\epsilon[\text{IoU}]$ and $\epsilon[\text{PE}]$, respectively.

To measure performance in terms of semantic labeling and precision we use polygon Average Precision, pAP, similarly defined as AP metrics for object detection. It is the area under the precision recall curve. A detection P_k with label q_k is a true positive if there is a ground truth polygon \hat{P}_m with label $y_m = q_k$ and $\text{IoU}(\hat{P}_m, P_k) < \gamma$. Only one detection may be matched to each ground truth polygon, extra detections are labeled false positives. For each label we calculate pAP^γ for $\gamma \in [0.5, 0.95]$ with step length 0.05. Then $\text{pAP}^m = \text{mean}(\text{pAP}^\gamma)$ and we average over labels to get mpAP^m .

5. Experiments

The models are trained and evaluated on the Structured3D dataset [33] which is a large-scale, photo-realistic

Table 1: AP Scores for the 3 different models on generated data and joint prediction.

Model	NMS	pAP ^m	pAP ^m	pAP ^m	mpAP ^m
		wall	floor	ceiling	
CSP	✓	14.5	15.8	20.3	16.8
CSP		4.6	4.2	5.6	4.8
HGC	✓	93.6	90.5	97.6	93.9
HGC		93.4	88.9	97.5	93.3
HGC +	✓	55.8	39.6	65.9	53.8
HGC +		55.1	40.4	65.6	53.7

+ Joint Wireframe Detection.

simulated dataset with 3D structure annotations. It consists of 3’500 scenes, with a total of 21’835 rooms and 196’515 frames and has a predefined training, validation and test data split. We use the existing annotations for the perspective images which include information about the planes such as polygon, 3D parameters and semantic label. We evaluate performance between the two proposed models on simulated wireframe detections and compare the HGC model with state of the art Room Layout Estimation when using both simulated wireframe detections and HAWP detections. The supplementary material also includes an experiment on inference time.

5.1. Semantic Plane Detection on Room Layout

In this experiment we use two different approaches to the Wireframe detector. The most frequently used is a synthetic detector where line segments are generated from the ground truth annotations. Given an image I_i with M_i line segments $\hat{\mathbf{L}}_i$ and junctions $\hat{\mathbf{J}}_i$ we generate line segments $\mathbf{L}_i^G = \{L_1^G, L_2^G, \dots, L_B^G\}$ according to L-CNN [34]. We generate $B = M_i$ for each image I_i . For the HGC model we also train a Wireframe detector jointly with the network. This is more challenging which is indicated by the footnote in Table 1. From the table we see that the HGC model is superior to the cycle sampling based model and is not as dependent on the NMS post-processing for performance. The HGC performance on synthetic detections shows the potential given a correctly tuned Wireframe detector. Performance declines when jointly predicting the Wireframe and the Polygons, likely due to worse wireframes and less flexibility in the model. It is also clear that the floor is the most difficult type of polygon, followed by wall and finally ceiling, likely due to occlusions.

5.2. Room Layout from Wireframe Detections

Furthermore, we compare the best HGC models from both synthetic and joint wireframe detections with state of the art models for Room Layout estimation. From Table 2 we see that the HGC model with synthetic detections out-

Table 2: Room Layout scores for different models on generated data and joint prediction compared with state of the art models.

Type	$\epsilon[\text{IoU}]$ (%)	$\epsilon[\text{PE}]$ (%)
RaC [20]	76.29	8.07
NonCuboid [27]	81.40	5.87
HGC	94.66	2.57
HGC +	50.76	41.27

+ Joint Wireframe Detection.

perform the state of the art models by a large margin for the 2D metrics $\epsilon[\text{IoU}]$ and $\epsilon[\text{PE}]$. However the HGC with joint wireframe detection cannot reach the same performance as competing models.

5.3. Qualitative Comparison of Models

Finally, we have a qualitative comparison between the models tested in this paper. In Figure 6 we see semantic polygon detections from synthetic wireframes. The first column shows the wireframe given as input to the model, the second column is the output of the Cycle Sampling Polygon-based classifier with NMS, the third column is the Heterogeneous Graph-based Classifier with NMS and finally column four holds the ground truth annotation.

In Figure 7 we see similar images but for the HGC model with joint Wireframe detection and NMS. Here images are alternating in pairs with the first being the detected wireframe and the second being the polygon output of the model.

6. Conclusions

In this paper we introduce a new neural network based system for semantic plane detection. The system takes a single image as input and outputs junctions, lines and planes corresponding to room layout estimation. The method uses an end-to-end learnt feature space, that is used both for detection and feature representations of junctions, lines and planes. These are then represented as a heterogeneous graph. This representation is refined using two heterogeneous graph transformers followed by a classification network to provide the final output. The methods are evaluated on the Structured3D dataset. The experiments show that given sufficiently good wireframe detections the model can outperform the state of the art Room Layout estimation models on 2D metrics and is a promising direction for research. Ideas for further studies could be *e.g.* how a polygon should be sampled for features or other ways of extracting polygons given edge scores. Another research topic would be to see if the model can add missing line segments in the wireframe using edge prediction on the graph.



Figure 6: Example output from the CSP and HGC model on simulated wireframe detections. From the left (i) Ground truth with sampled wireframe on. (ii) CSP estimate. (iii) HGC estimate. (iv) Ground truth polygons.



Figure 7: Example output from the HGC model when jointly detecting wireframes. Images are alternating (i) Ground truth with wireframe detections on, then (ii) HGC estimate.

References

- [1] James M. Coughlan and Alan L. Yuille. Manhattan world: compass direction from a single image by bayesian inference. In *Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, pages 941–947 vol.2, 1999. 2
- [2] Luca Del Pero, Joshua Bowdish, Daniel Fried, Bonnie Kermgard, Emily Hartley, and Kobus Barnard. Bayesian geometric modeling of indoor scenes. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2012*, Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 2719–2726, Oct. 2012. 2012 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2012 ; Conference date: 16-06-2012 Through 21-06-2012. 2
- [3] David Gillsjö, Gabrielle Flood, and Kalle Åström. Semantic room wireframe detection from a single view. In *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 1886–1893, 2022. 2, 3
- [4] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017. 3
- [5] Varsha Hedau, Derek Hoiem, and David Forsyth. Recovering the spatial layout of cluttered rooms. In *2009 IEEE 12th international conference on computer vision*, pages 1849–1856. IEEE, 2009. 2
- [6] Henry Howard-Jenkins, Shuda Li, and Victor Adrian Prisacariu. Thinking outside the box: Generation of unconstrained 3d room layouts. In *ACCV*, 2018. 2
- [7] Ziniu Hu, Yuxiao Dong, Kuansan Wang, and Yizhou Sun. Heterogeneous graph transformer. In *Proceedings of The Web Conference 2020, WWW '20*, page 2704–2710, New York, NY, USA, 2020. Association for Computing Machinery. 3, 4
- [8] Ivan Kruzhilov, Mikhail Romanov, Dmitriy Babichev, and Anton Konushin. Double refinement network for room layout estimation. In *Pattern Recognition*, pages 557–568, Cham, 2020. Springer International Publishing. 2
- [9] Boon Chai Lee. Algorithmic approaches to circuit enumeration problems and applications. Technical report, Cambridge, Mass.: Massachusetts Institute of Technology, Dept. of ..., 1982. 3
- [10] Chen-Yu Lee, Vijay Badrinarayanan, Tomasz Malisiewicz, and Andrew Rabinovich. Roomnet: End-to-end room layout estimation, 2017. 2
- [11] Hung Jin Lin, Sheng-Wei Huang, Shang-Hong Lai, and Chen-Kuo Chiang. Indoor scene layout estimation from a single image. In *2018 24th International Conference on Pattern Recognition (ICPR)*, 2018. 2
- [12] Hung Jin Lin and Shang-Hong Lai. Deeproom: 3d room layout and pose estimation from a single image. In Shivakumara Palaiahnakote, Gabriella Sanniti di Baja, Liang Wang, and Wei Qi Yan, editors, *Pattern Recognition*, pages 719–733, Cham, 2020. Springer International Publishing. 2
- [13] Chen Liu, Kihwan Kim, Jinwei Gu, Yasutaka Furukawa, and Jan Kautz. Planercnn: 3d plane detection and reconstruction from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [14] C. Liu, J. Yang, D. Ceylan, E. Yumer, and Y. Furukawa. Planenet: Piece-wise planar reconstruction from a single rgb image. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2579–2588. IEEE Computer Society, 2018. 2
- [15] Arun Mallya and Svetlana Lazebnik. Learning informative edge maps for indoor scene layout prediction. In *2015 International Conference on Computer Vision, ICCV 2015*, Proceedings of the IEEE International Conference on Computer Vision, pages 936–944, United States, Feb. 2015. Institute of Electrical and Electronics Engineers Inc. 15th IEEE International Conference on Computer Vision, ICCV 2015 ; Conference date: 11-12-2015 Through 18-12-2015. 2
- [16] Quan Meng, Jiakai Zhang, Qiang Hu, Xuming He, and Jingyi Yu. *LGNN: A Context-Aware Line Segment Detector*, page 4364–4372. Association for Computing Machinery, New York, NY, USA, 2020. 2
- [17] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *Computer Vision - 14th European Conference, ECCV 2016, Proceedings*, Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), pages 483–499, Germany, 2016. Springer Verlag. 3
- [18] Yuzhuo Ren, Shangwen Li, Chen Chen, and C.-C. Jay Kuo. A coarse-to-fine indoor layout estimation (cfle) method. In Shang-Hong Lai, Vincent Lepetit, Ko Nishino, and Yoichi Sato, editors, *Computer Vision – ACCV 2016*, pages 36–51, Cham, 2017. Springer International Publishing. 2
- [19] Alexander G. Schwing, Tamir Hazan, Marc Pollefeys, and Raquel Urtasun. Efficient structured prediction for 3d indoor scene understanding. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2012*, Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 2815–2822, 2012. 2
- [20] Sinisa Stekovic, Shreyas Hampali, Mahdi Rad, Sayan Deb Sarkar, Friedrich Fraundorfer, and Vincent Lepetit. General 3d room layout from a single view by render-and-compare. In *Computer Vision – ECCV 2020*, pages 187–203, Cham, 2020. Springer International Publishing. 2, 6, 7
- [21] Bin Tan, Nan Xue, Song Bai, Tianfu Wu, and Gui-Song Xia. Planetr: Structure-guided transformers for 3d plane recovery. In *International Conference on Computer Vision*, 2021. 2
- [22] Gui-Song Xia, Julie Delon, and Yann Gousseau. Accurate junction detection and characterization in natural images. *International journal of computer vision*, 106(1):31–56, 2014. 2
- [23] Nan Xue, Song Bai, Fudong Wang, Gui-Song Xia, Tianfu Wu, and Liangpei Zhang. Learning attraction field representation for robust line segment detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1595–1603, 2019. 2
- [24] Nan Xue, Song Bai, Fu-Dong Wang, Gui-Song Xia, Tianfu Wu, Liangpei Zhang, and Philip HS Torr. Learning regional attraction for line segment detection. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 2

- [25] Nan Xue, Tianfu Wu, Song Bai, Fu-Dong Wang, Gui-Song Xia, Liangpei Zhang, and Philip H.S. Torr. Holistically-attracted wireframe parsing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3
- [26] Chenggang Yan, Biyao Shao, Hao Zhao, Ruixin Ning, Yongdong Zhang, and Feng Xu. 3d room layout estimation from a single rgb image. *IEEE Transactions on Multimedia*, 22(11):3014–3024, 2020. 2
- [27] Cheng Yang, Jia Zheng, Xili Dai, Rui Tang, Yi Ma, and Xiaojun Yuan. Learning to reconstruct 3d non-cuboid room layout from a single rgb image. In *WACV*, 2022. 2, 7
- [28] Fengting Yang and Zihan Zhou. Recovering 3d planes from a single image via convolutional neural networks. In Vittorio Ferrari, Martial Hebert, Cristian Sminchisescu, and Yair Weiss, editors, *Computer Vision – ECCV 2018*, pages 87–103. Springer International Publishing, 2018. 2
- [29] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop, 2016. 2
- [30] Zehao Yu, Jia Zheng, Dongze Lian, Zihan Zhou, and Shenghua Gao. Single-image piece-wise planar 3d reconstruction via associative embedding. In *CVPR*, pages 1029–1037, 2019. 2
- [31] Weidong Zhang, Wei Zhang, and Jason Gu. Edge-semantic learning strategy for layout estimation in indoor environment. *CoRR*, abs/1901.00621, 2019. 2
- [32] Hao Zhao, Ming Lu, Anbang Yao, Yiwen Guo, Yurong Chen, and Li Zhang. Physics inspired optimization on semantic transfer features: An alternative method for room layout estimation, 2017. 2
- [33] Jia Zheng, Junfei Zhang, Jing Li, Rui Tang, Shenghua Gao, and Zihan Zhou. Structured3d: A large photo-realistic dataset for structured 3d modeling. In *Proceedings of The European Conference on Computer Vision (ECCV)*, 2020. 2, 6
- [34] Yichao Zhou, Haozhi Qi, and Yi Ma. End-to-end wireframe parsing. In *ICCV 2019*, 2019. 7