

# Holistic Geometric Feature Learning for Structured Reconstruction

Ziqiong Lu\*    Linxi Huan\*    Qiyuan Ma    Xianwei Zheng†  
The State Key Lab. LIESMARS, Wuhan University  
{zql2018, whu\_hlx, qiyuanma, zhengxw}@whu.edu.cn

## Abstract

*The inference of topological principles is a key problem in structured reconstruction. We observe that wrongly predicted topological relationships are often incurred by the lack of holistic geometry clues in low-level features. Inspired by the fact that massive signals can be compactly described with frequency analysis, we experimentally explore the efficiency and tendency of learning structure geometry in the frequency domain. Accordingly, we propose a frequency-domain feature learning strategy (F-Learn) to fuse scattered geometric fragments holistically for topology-intact structure reasoning. Benefiting from the parsimonious design, the F-Learn strategy can be easily deployed into a deep reconstructor with a lightweight model modification. Experiments demonstrate that the F-Learn strategy can effectively introduce structure awareness into geometric primitive detection and topology inference, bringing significant performance improvement to final structured reconstruction. Code and pre-trained models are available at <https://github.com/Geo-Tell/F-Learn>.*

## 1. Introduction

The structured reconstruction models the shape grammar/topology that depicts procedural shape generation [24], which can facilitate various downstream vision tasks, such as feature matching [19, 44], 3D modeling [23, 32] and 3D scene understanding [47, 42]. The recovery of a shape topology is generally realized via two steps: (i) geometric primitive detection and (ii) graph inference based on extracted geometric primitives (e.g., corners and edges). Following this primitive-to-structure principle, previous researchers usually resort to heatmaps of corners and edges for subsequent structure reasoning [18, 47, 42, 2]. With low-level parsing tasks promoted by deep learning techniques, many works have greatly improved the performance of high-level structure recovery in recent years [42, 24, 2].

Researchers commonly dedicate to graph inference with geometric features generated by proven hierarchical backbones (e.g., ResNet [10]) for structured reconstruction. However, incorrect topology inference remains a significant problem due to the lack of holistic geometric clues in the low-level features. Specifically, the geometric fragments extracted by shallow layers can hardly be fused into structurally informative features, leading to wrongly reasoned topological principles (as shown in Fig. 1 (a)). Different from the prior arts that focus on graph inference, this paper turns to study an efficient strategy for holistically learning structure-related features with low-level geometric fragments.

Low-level geometric features fundamentally support high-level feature extraction and structure recovery. Specifically, the low-level shallow-layer feature maps provide information for precise geometry localization, which is essential for high-level primitive detection and topology inference. Restricted by the limited receptive field, it is hard to holistically capture structure geometries in the shallow layers. We consider that the low efficiency of processing low-level geometric fragments in the space domain is the root cause.

Inspired by the fact that low-level features can be compactly encoded in the frequency domain, we are interested in achieving efficient geometry feature learning with frequency analysis. In the frequency map of a given feature map, every value encodes the global information of the feature map regarding the corresponding frequency. Therefore, performing convolution operations in the frequency domain actually means directly combining the spatial information in a holistic manner.

To be concrete, given a set of low-level feature maps  $\{f_n(x, y)\}$  and the corresponding frequency map  $\{F_n(u, v)\}$ , the information encoded at  $(u_i, v_j)$  describes a certain changing pattern of signals in the space domain. When applying a convolution operation to a frequency position  $(u_i, v_j)$ , the feature components of related frequencies in every map of  $\{f_n(x, y)\}$  will be merged across the channel dimension in the space domain. As the low-level geometric information generally belongs to high-frequency sig-

\*Equal contribution.

†Corresponding author.

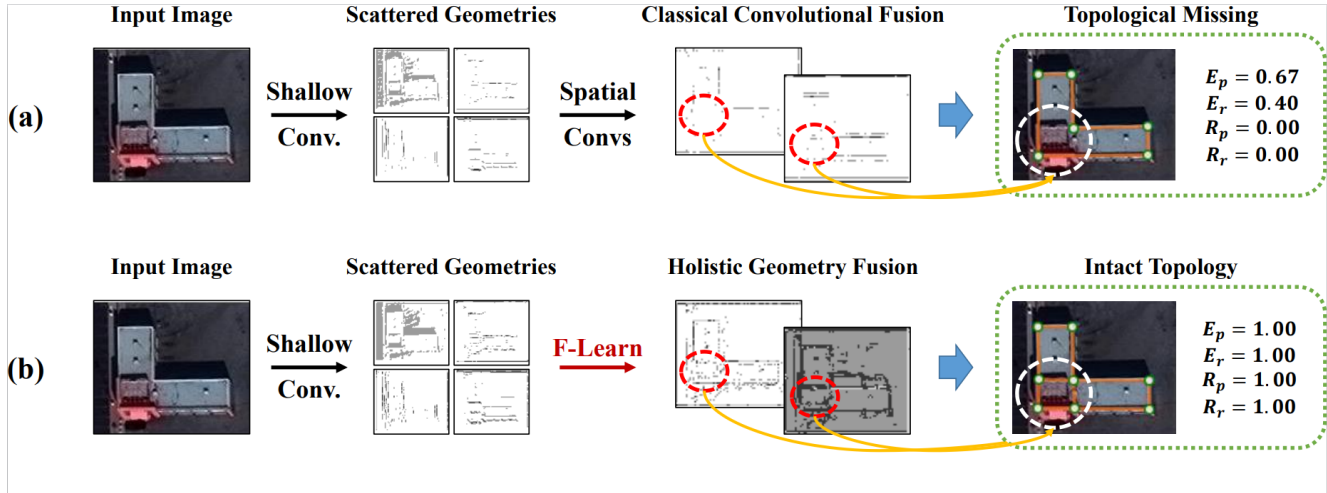


Figure 1. Comparison of learning low-level features in the space and frequency domains. (a) Structure reconstruction with low-level geometric features learned **in the space domain**. The inefficient fusion of low-level geometric fragments loses the structure clues of the bottom left roof region and consequently results in missing topological relations in the final reconstructed structure graph. (b) Structure reconstruction with low-level geometric features learned **with the proposed frequency-domain feature learning strategy (F-Learn)**. With the geometric features compactly processed in the frequency domain, our F-Learn strategy effectively achieves holistic geometry fusion for inferring the right topological principle for structure reconstruction.  $E_p$  and  $E_r$  denote the scores of precision and recall for detected edges, while  $R_p$  and  $R_r$  refer to those of reconstructed regions.

nals, a frequency-domain convolution directly realizes the integration of geometric primitives scattered in  $\{f_n(x, y)\}$ . The frequency-domain convolution thereby works more efficiently in holistic geometry learning than the space-domain counterpart, which combines local geometric clues without a global view.

Based on the discussion above, we propose a frequency-domain feature learning strategy (F-Learn) to efficiently extract holistic geometry features for guiding the inference of structure topology. Working in the frequency domain, the F-Learn strategy holistically fuses the separated geometric primitives into structurally informative features. The proposed F-Learn strategy can be readily applied to a primitive-to-structure framework for structured reconstruction.

In summary, the main contributions of this paper are:

- Exploration with a simple geometry recovery task sheds light on the difference in learning tendency between frequency- and space-domain convolutions. The results also validate the high efficiency of frequency-domain convolution in learning holistic geometry.
- A parsimonious frequency-domain feature learning strategy (F-Learn) is proposed to generate structurally informative geometry features for structured reconstruction.
- Experiments on vectorizing world buildings demonstrate that our F-Learn strategy greatly improves the performance of structured reconstruction.

## 2. Related Work

### 2.1. Geometric Primitive Detection

The detection of geometric primitives is a long-standing vision task that has been extensively explored with hand-crafted descriptors in the early stage [26, 9]. Recently, with the learning ability of hierarchical features, deep learning models have significantly promoted the completeness and precision of detected geometric primitives. A popular way for low-level geometry extraction is to perform heatmap regression or pixel-wise binary classification. The heatmap regression technique is widely adopted in corner detection. For instance, methods like CornerNet [17] and CenterNet [6] learn corner heatmaps for the downstream task of object detection. Binary classification is often used to localize edge pixels. Representative edge detectors like HED [34] and RCF [21] focus on fusing multi-level features for edge pixel classification, while some works pay attention to learning crisp edges or semantic contours [40, 5, 11].

Compared to the pixel-wise inference of corners and edges, the detection of line primitives is more structural as a line segment should be defined by two endpoints [7, 16]. Since the task of wireframe parsing was introduced by [14], increasing interest has been witnessed in inferring line segment candidates with learned junctions [45, 38, 20]. With the advance in low- and mid-level primitive detection, researchers have started to make efforts on structured reconstruction in recent years.

## 2.2. Structured Reconstruction

The structured reconstruction requires reasoning the overall topology of a given shape. Generally, structured reconstruction is a large research field that contains various tasks ranging from 3D object CAD modeling [32], layout estimation [18, 47], to roof extraction [43]. For instance-level structure reconstruction, many methods resort to generative models [32] or detecting key points with fixed topologies (e.g., skeletons of human bodies and hands) [33, 46].

As for scene-level structural modeling, the reconstruction targets usually belong to special semantic-related structures, such as planar building roofs and indoor floorplans. For these tasks, the derived structures have to match semantic regions with robustness to other irrelevant structure information. For example, room layout estimation has to ensure that connected lines can form wall regions without extra lines from windows and doors. Related approaches can be roughly categorized into two kinds: (i) the local primitive based and (ii) the global information based. Pioneering local primitive based works derive structural representation from an image by post-processing heatmaps of corners and edges [18, 47].

With respect to the global information based methods, it is a popular choice to form the final structure graph with structure-related regions segmented from the input image [41, 28]. Different from the segmentation based methods, recent studies consider a holistic graph inference with geometric primitives as nodes for high-level planer extraction. For example, Conv-MPN [42] adopts a variant of the graph neural network to pass messages across the whole graph, and HEAT [2] learns the topological pattern of edges with an attention transformer. Beyond exploring the high-level information in structure recovery, we are interested in improving the efficiency of exploiting low-level geometric features, which serve as the foundation of accurate structure topology inference.

## 2.3. Frequency Analysis in Deep Learning

Frequency analysis is a technique that provides a compressed representation of signals in the frequency domain. Not only commonly used in classical digital signal processing, but frequency analysis also remains powerful in the era of deep learning. [37] and [1] investigate the training and generalization of neural networks via frequency analysis. [35] analyzes the spectral bias of deep models regarding multiple vision parsing tasks. [8] applies frequency analysis to generated images for improving image synthesis quality, while [15] introduces a focal frequency loss that forces generative models to learn hard frequencies. [25] designs a frequency channel attention mechanism to compress channel-wise information with scalars. [3] and [31] utilize frequency analysis in model compression. More applications of the

frequency-domain representation can also be found in domain adaptation [12, 13] and position embedding [22, 29].

Although frequency analysis has been combined with deep learning techniques in recent years, it is generally used as a mature tool in prior arts without a deep investigation of the reasons why frequency analysis is useful. To this end, we experimentally studied the behavior tendency of frequency-domain convolutions in learning holistic geometries. Based on the investigation results, we designed the F-Learn strategy to efficiently learn holistic geometry clues in low-level feature maps for structured reconstruction.

## 3. Method

Geometric features extracted at the early convolutional stage fundamentally support the structure reconstruction with precise geometry localization clues. Therefore, we are interested in addressing a key problem for learning reliable low-level geometric features: the low efficiency of fusing geometric fragments in shallow convolution layers. Considering that geometric information is generally high-frequency signals, we are motivated to study the problem mentioned before with frequency analysis for structure reconstruction. In the following of this section, we first present the preliminaries of frequency analysis. Next, we validate the high efficiency of learning in the frequency domain for holistic geometry recovery. Finally, we introduce our frequency-domain feature learning strategy (F-Learn) and the application of the F-Learn in a given structure reconstruction model.

### 3.1. Preliminaries of Frequency Analysis

The discrete Fourier transform (DFT) is necessary for analyzing image data in the frequency domain. With  $I(x, y)$  denoting the color signal at the spatial position  $(x, y)$  of an  $M \times N$  image, the DFT converts the information of  $I$  to a frequency map  $F$  in the frequency domain by Eq. (1).

$$F(u, v) = \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} I(x, y) e^{-j2\pi(ux/M+vy/N)}. \quad (1)$$

In Eq. (1),  $F(u, v)$  is a value that encodes a two-dimension signal changing that belongs to the frequency determined by  $u$  and  $v$ . The frequency map excels in holistically representing signals that show a similar frequency pattern. Therefore, we assume that the DFT is a powerful tool for efficiently learning structure-related geometry features. To validate this hypothesis, we conducted experiments with a geometry recovery task, and analyze the results quantitatively and qualitatively as below.

### 3.2. Learning Efficiency Analysis

In this section, we study the learning efficiency of fusing geometric fragments in the space and frequency domains

with a simple geometry recovery task. We first decomposed a geometric structure composed of a circle and a square into several geometric fragments as illustrated in Fig. 2 (a), and left overlapping between neighboring fragments to avoid trivial solutions.

With the purpose to recover the binary image of the original structure with the geometric fragments, we studied the learning efficiency of frequency-domain and space-domain convolution models. We constructed a frequency-domain learning model (F-Learn) and several space-domain convolution models. The baseline was set as a space-domain convolution module (BConv) built with two  $1 \times 1$  convolution layers and one  $3 \times 3$  layer in between. The F-Learn model consists of three key components: (1) a DFT; (2) two parallel convolution modules same with the baseline for separately processing the real and imaginary parts; and (3) an inverse DFT (IDFT) followed by magnitude computation of complex values. For a fair comparison, we also combined a pair of baseline modules into two kinds of space-domain learning models with parallel and cascade structures, namely Conv-Casv1 and Conv-Parv1. Extra  $3 \times 3$  convolutions were also added into the cascade and parallel space-domain models to simulate the function of DFT and IDFT. The modified models were named Conv-Casv2 and Conv-Parv2 for simplicity. For all models, the channel number of the intermediate feature maps was 64, and a  $1 \times 1$  convolution was used as the final binary classifier. Detailed model settings are depicted in Fig. 2(b).

The experiments were conducted on  $128 \times 128$  binary images of the geometric structure and fragments. The training of all models was completed with an Adam optimizer. The learning rate and training period were set to 0.1 and 50 epochs, respectively. The binary cross entropy loss (BCELoss) was chosen as the training guide. To alleviate the influence of randomness, we generated 100 seeds to set 100 trials, and all models were trained with the same seed in every trial. The performance of all models was evaluated with an F1-score measurement at a threshold of 0.5. Quantitative and qualitative results are shown in Fig. 3 and Fig. 4.

Fig. 3 draws the averaged F1-score curves of 100 trials. Compared with the space-domain counterparts, the frequency-domain model (F-Learn) shows superiority in recovering the complete structure in a more efficient manner.

The visualized results in Fig. 4 also show that the frequency-domain learning model exceeds the basic spatial convolution model with faster convergence speed and better structure recovery quality.

Fig. 4 reveals that the F-Learn model easily rebuilds the original structure, while the space-domain methods encounter difficulties in efficiently reconstructing the circle. We also find the difference in learning tendency between the F-Learn and space-domain models. The F-Learn model pri-

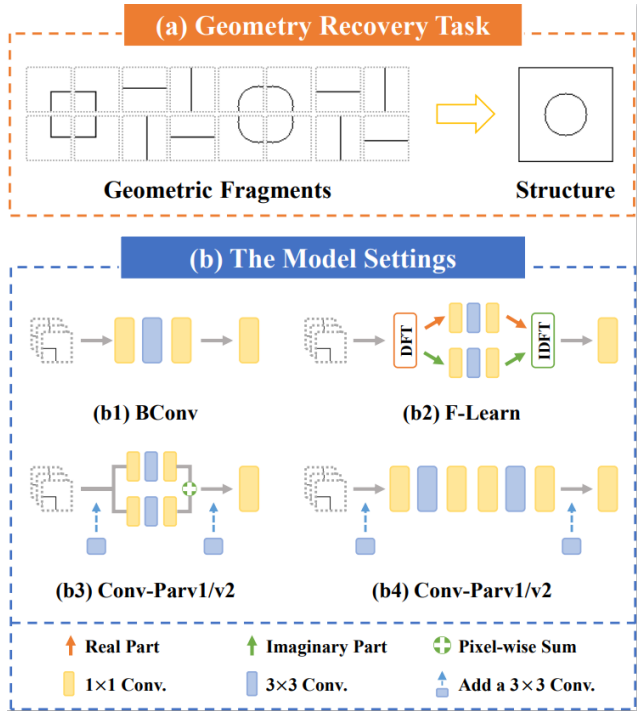


Figure 2. The settings of the simple geometry recovery task. (a) The task of recovering the holistic structure with geometric fragments. (b) The detailed settings of compared models.




Figure 3. The averaged F1-score curves of 100 trials for the F-Learn model and the compared space-domain methods.

oritizes holistic integration before further enlarging the gap between the foreground and background. In contrast, the space-domain models first generate high responses to local areas that have strong signal intensity in the input. Hence, the space-domain models prefer overlapped areas without considering structural continuity. In practice, the holistic learning pattern of F-Learn will benefit the topology inference in two ways: (i) holistic geometry clues are ready for learning high-level features at the early stage, which **alleviates the burden on learning structure-related low-level features with far supervision signal**; (ii) the holistic geometry learning excels in **preserving the connection in-**




Figure 4. The visualized results of F-Learn, Bconv and Cov-Parv1 at different epochs.

**formation** in the high-resolution low-level feature maps, which are often reused to **provide accurate structure localization** for topology inference.

We attribute the advantages of the frequency domain based model to its ability to holistically process signals that belong to similar frequencies. Based on the findings above, we designed a frequency-domain feature learning strategy to generate holistic geometric clues in the low-level feature maps for structured reconstruction.

### 3.3. Frequency-domain Learning Strategy

The proposed frequency-domain learning strategy (F-Learn) inherits the architecture shown in Fig. 2. As the F-Learn strategy will face more complex data and larger batch size in practice than in Sec. 3.2, we replace each convolution layer between DFT and IDFT with a combination (CBR) of a convolution layer, a batch normalization operation, and a relu activation function. We refer to all operations between DFT and IDFT collectively as a CBR group (CBRG). Given a set of geometric feature maps  $\{f_i | i = 0, \dots, N\}$ , the F-Learn strategy is formulated as

$$\begin{aligned} (F_i^{\text{re}}, F_i^{\text{im}}) &= \text{DFT}(f_i); \\ \{\hat{F}_i^{\text{re}}, \hat{F}_i^{\text{im}}\} &= \text{CBRG}(\{F_i^{\text{re}}\}), \text{CBRG}(\{F_i^{\text{im}}\}); \\ \{\hat{f}_i\} &= |\text{IDFT}(\{\hat{F}_i^{\text{re}}, \hat{F}_i^{\text{im}}\})|. \end{aligned} \quad (2)$$

The F-learn strategy first converts every  $f_i$  into a pair of maps that record the real and imaginary parts of the DFT, i.e.,  $F_i^{\text{re}}$  and  $F_i^{\text{im}}$ . The maps of the real part  $\{F_i^{\text{re}}\}$  are subsequently processed by a CBRG that contains two  $1 \times 1$  layers with a  $3 \times 3$  one in between, and the same operation is also applied to  $\{F_i^{\text{im}}\}$ . In the frequency domain, the  $1 \times 1$  convolution directly fuses the information of the same frequency across the channel dimension, while the  $3 \times 3$  counterpart enhances features with signals that have similar frequencies. The high-frequency geometric fragments

in  $\{f_i\}$  are easily combined and enhanced in such a holistic way. The structurally informative space-domain features  $\{\hat{f}_i\}$  can be finally obtained through an IDFT along with magnitude computation  $|\cdot|$ . With the parsimonious design, the F-learn strategy is readily inserted into a hierarchical convolutional backbone for holistic geometry feature learning.

For feature propagation, the  $\{f_i\}$  and  $\{\hat{f}_i\}$  are fused into  $\{\tilde{f}_i\}$  via

$$\{\tilde{f}_i\} = C_{1 \times 1} \text{BR}(\text{Concat}(\{f_i\}, \{\hat{f}_i\})), \quad (3)$$

where Concat and  $C_{1 \times 1} \text{BR}$  denote the concatenation operation and a  $1 \times 1$  convolution based CBR, respectively.

The enhanced geometric features  $\{\tilde{f}_i\}$  are then fed into the next convolution stage for higher-level feature learning and used to offer precise geometric localization for later structure inference.

The F-Learn strategy can be easily inserted into a convolution base backbone. With the commonly-used ResNet50 backbone [10] as an example, the F-Learn strategy can be directly deployed after the first convolution to holistically learn geometric features. In experiments, we implemented our F-Learn strategy in a state-of-the-art approach, i.e., the holistic edge attention transformer (HEAT) [2], for structured reconstruction.

## 4. Experiments

### 4.1. Experiment Settings

#### 4.1.1 Dataset

We tested our method on a dataset introduced by [24] for vectorizing world buildings. This dataset is built on the SpaceNet dataset [30] with 2001 images annotated with roof planar graphs. Each image contains a building instance, and the image is processed into a size of  $256 \times 256$ . In line with the prior arts [2], the training/validation/testing split is set to 1601/50/350.

#### 4.1.2 Implementation Details

We evaluated the F-learn strategy with the HEAT model for structured reconstruction. The HEAT model was built on a ResNet50 backbone pretrained on ImageNet [4]. We performed our experiments with Pytorch [27] in Python3.7 and used a workstation with one NVIDIA RTX 3090 GPU. We adopted the same settings of loss functions as the HEAT model. The training of the model was completed with an AdamW optimizer, and the training period lasts 800 epochs. The learning rate was initialized as  $2e-4$  and multiplied by a factor of 0.1 for the last 25% epoch. In line with previous works [2, 24], we used precision, recall, and F1 score to evaluate the quality of structure reconstruction in terms of corner detection, edge inference, and region recovery.

Table 1. **Quantitative comparison between the F-Learn strategy and other methods in terms of corners detection, edge extraction, and region reconstruction.** Prec and F1 are the abbreviations of the precision and f1-score metrics. The higher the scores are, the better the performance is. The best results are marked **bold**. (unit:%)

Evaluation Type →			Corner			Edge			Region		
	Method	Fully-neural	Joint	Prec	Recall	F1	Prec	Recall	F1	Prec	Recall
IP [24]	-	-	-	-	74.5	-	-	53.1	-	-	55.7
Exp-Cls [43]	-	-	92.2	75.9	83.2	75.4	60.4	67.1	74.9	54.7	63.5
ConvMPN [42]	✓	-	78.0	79.7	78.8	57.0	59.7	58.1	52.4	56.5	54.4
HAWP [39]	✓	✓	90.9	81.2	85.7	76.6	68.1	72.1	74.1	55.4	63.4
LETR [36]	✓	✓	87.8	74.8	80.8	59.7	58.6	59.1	68.3	48.7	56.8
HEAT [2]	✓	✓	91.7	83.0	87.1	80.6	72.3	76.2	76.4	65.6	70.6
HEAT(retrain) [2]	✓	✓	91.6	83.0	87.1	80.3	72.4	76.1	75.5	65.3	70.0
<b>F-Learn (Ours)</b>	✓	✓	<b>93.2</b>	<b>84.4</b>	<b>88.6</b>	<b>83.6</b>	<b>75.0</b>	<b>79.1</b>	<b>79.5</b>	<b>68.1</b>	<b>73.4</b>

### 4.1.3 Competing Methods

We compare the F-Learn strategy with six methods: HEAT [2], ConvMPN [42], IP [24], Exp-clis [43], HAWP [39] and LETR [36].

**HEAT** is an attention-based method that takes a 2D raster image as an input and reconstructs a planar graph in an end-to-end manner. HEAT works via three steps: 1) extracting hierarchical features by ResNet50; 2) detecting corners with multi-scale features enriched by a deformable attention module; and 3) employing two weight-sharing Transformer decoders to classify edges and reason structures with detected corners. Our F-Learn strategy is adopted in the first step for learning structure-related features.

Conv-MPN uses a graph neural network to infer edges with a pre-trained corner detector. IP firstly detects geometric primitives and then reconstructs a planar graph through several post-processing steps. Exp-clis is based on the geometric primitives produced by other methods (e.g., IP and Conv-MPN) and reconstructs a planar graph through an explore-and-classify framework. LCNN and HAWP are methods specifically proposed for wireframe parsing. LETR is a transformer-based method that directly generates lines without post-processing and heuristic guidance.

## 4.2. Comparison and Analysis

### 4.2.1 Quantitative Analysis

The quantitative results displayed in Tab. 1 show that our F-Learn strategy advances roof structure reconstruction with state-of-the-art performance. The F-Learn strategy edges out other compared methods by significant performance gains regarding corner detection, edge inference, and region reconstruction under all metrics.

With respect to the detection of corners and edges, the F-Learn strategy outperforms the second-best by at least 1.0% and 3.0% in the precision measurement, while 1.4% and

2.6% under the recall metric. The improvement in extracting low-level geometric primitives indicates the outstanding ability of the F-Learn strategy for providing holistic and accurate geometric clues. The F-Learn strategy also brings at least a 2.8% gain in F1 score in terms of region reconstruction, which reveals the assistance of the F-Learn strategy in inferring correct topological relations. Besides, it is worth noticing that the F-Learn strategy significantly outperforms the original HEAT model under all metrics, but only brings a little computation increase compared to HEAT. Tab. 2 presents a detailed comparison of computation requirements.

Table 2. **Quantitative evaluation of the change of computing efficiency brought by F-Learn.** Param. and FLOPs are the abbreviations of the parameter quantity and floating point operations, used to measure the complexity of an algorithm/model. Time per Image represents the time it takes the model to infer an image and is used to evaluate the calculation speed of the model.

	Param.	FLOPs	Time per Image
w/o F-Learn	68.747M	670.919G	0.10s
w/ F-Learn	68.847M	684.039G	0.11s

In Tab. 2, the computation increase of F-Learn strategy in terms of parameters, FLOPs, and inference time per image can be clearly seen, where the amount is little. Especially, the parameter increase brought by F-Learn strategy is only 0.145%, which is almost negligible. The results further demonstrates that the F-Learn strategy is simple yet effective.

### 4.2.2 Qualitative Analysis

For perceptual comparison, we also present the visualized reconstruction results in Fig. 5. Compared to other methods, our F-Learn is capable of detecting low-level geometric

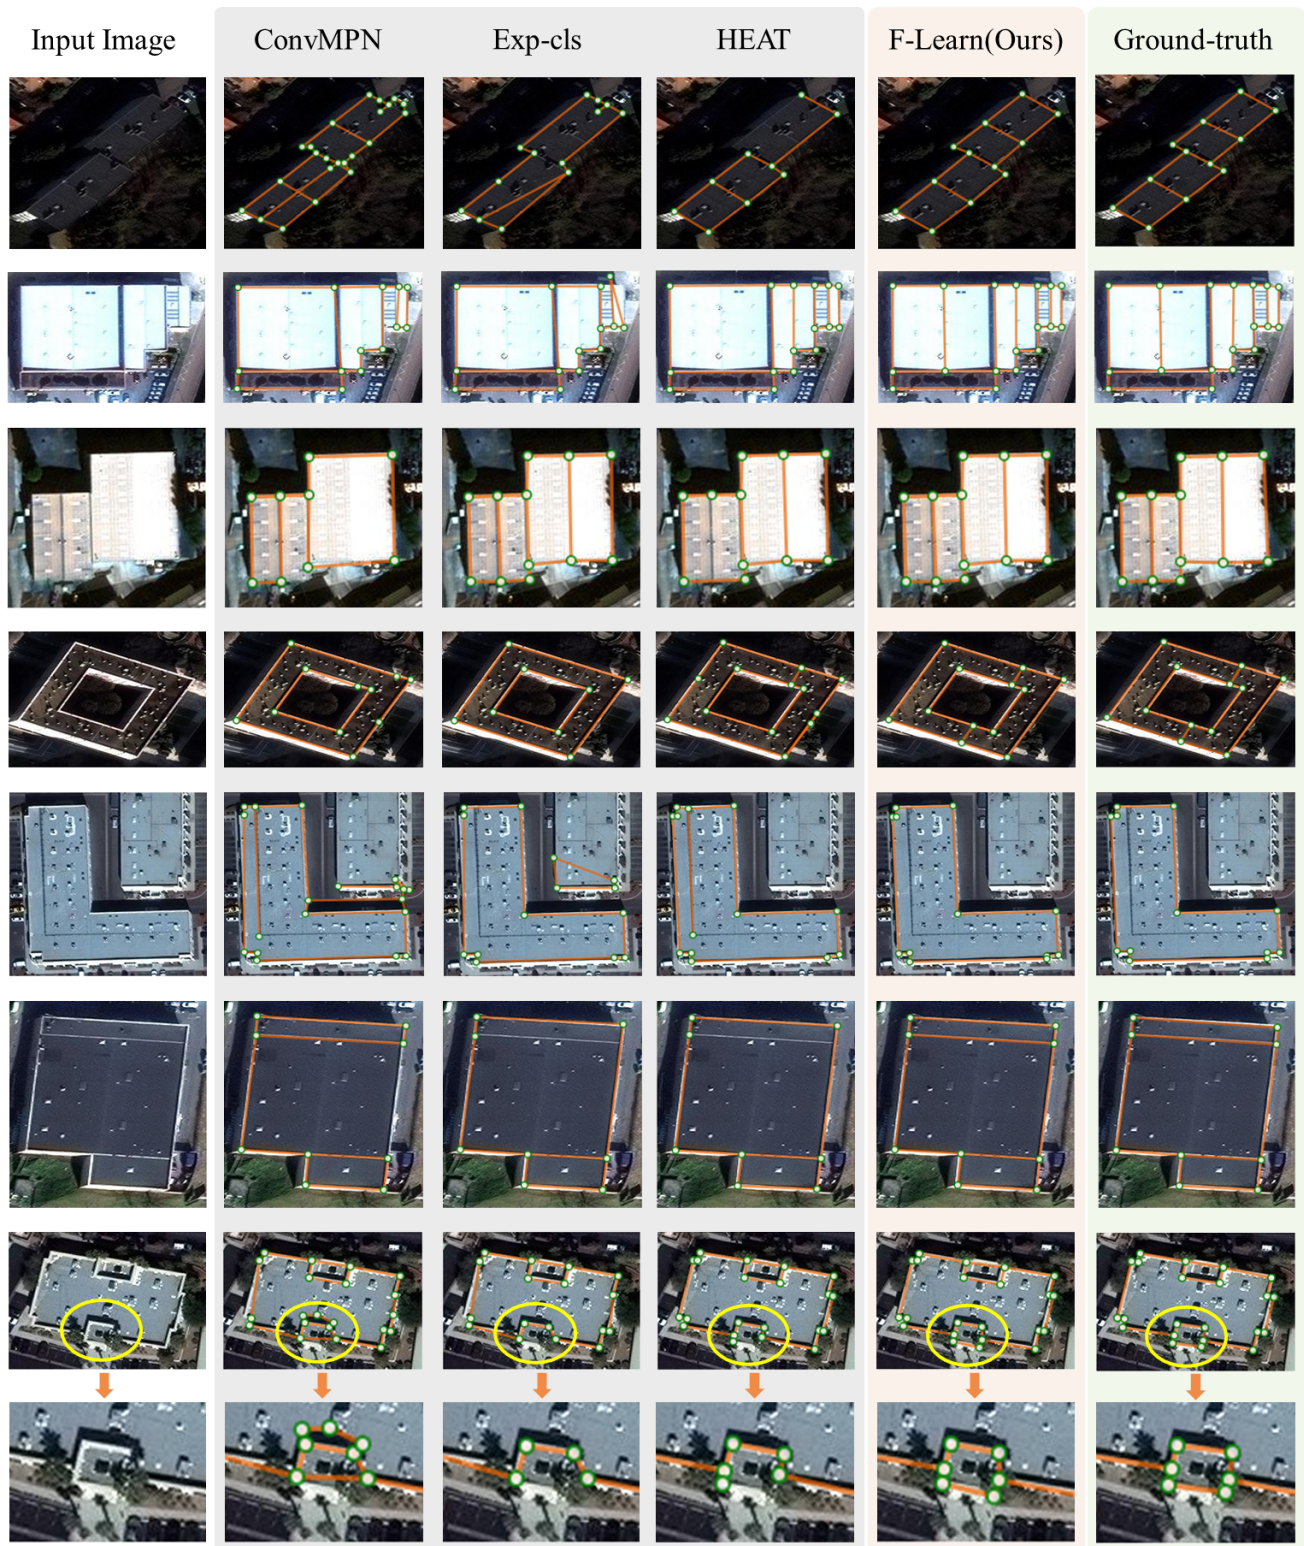


Figure 5. **Qualitative results on outdoor roof structure reconstruction.** Other methods build wrong topological relations due to redundant or missing geometric primitives. Our F-Learn strategy derives the correct topological principles with corners and edges more fitted to the holistic structures.

primitives that are consistent with the holistic structure.

When reconstructing the roof shown in the first row, ConvMPN and Exp-cls tend to detect redundant corners due to the shadow inference and therefore wrongly separate the roof regions. The original HEAT outputs a non-closed roof structure because key edges are missing. In contrast, the corners and edges learned with the F-Learn strategy are better fitted into the overall structure. This is because the F-Learn strategy can offer low-level features that are rich in holistic geometric clues for topology-preserved roof structure reconstruction. The results in the following rows additionally validate that the F-Learn strategy can effectively extract geometric primitives that are easily ignored by other methods, especially the original HEAT model. By comparing the inferred planar graphs in the fifth row, one can also find that the F-Learn improves the HEAT model with robustness to irrelevant structure information.

### 4.3. Ablation Studies

#### 4.3.1 Comparison with Space-domain Methods

In addition to the exploration with a simple geometry fusion task in Sec. 3.2, we further study the F-Learn strategy with the space-domain learning strategies in the real scene for roof structure reconstruction. Similar to Sec. 3.2, we construct a space-domain learning baseline (BConv) with a basic module composed of two  $1 \times 1$  convolution layers and one  $3 \times 3$  layer in between. Furthermore, we design two more space-domain strategies with the basic modules arranged in parallel and cascade, denoted as Conv-Parv. and Conv-Casv., respectively. Same with the F-Learn, additional operations of batch normalization and ReLU activation are added after each convolution layer used in every space-domain strategy.

Table 3. **Quantitative comparisons with space-domain learning methods.** The F-Learn strategy is compared with the space-domain counterparts under the F1-score measurement in terms of corner detection, edge inference, and region reconstruction. The higher the scores are, the better the performance is. The best results are marked **bold**. (unit:%)

Method	Corner	Edge	Region
HEAT	87.1	76.1	70.0
HEAT-BConv	87.4	76.2	70.7
HEAT-Conv-Cas.	88.1	77.1	70.2
HEAT-Conv-Par.	87.7	76.8	69.5
<b>F-Learn</b>	<b>88.6</b>	<b>79.1</b>	<b>73.4</b>

With the F1-score measurement, Tab. 3 presents the performance of the F-learn and space-domain learning strategies in detecting geometric primitives. Although some space-domain methods are comparable to the proposed F-

Learn strategy in corner extraction, it can be seen that the F-Learn strategy gains over these methods in detecting edges and regions by at least 2% and 2.7%, respectively. Because edges and regions are primitives closely related to roof topology, this phenomenon demonstrates the holistic geometry learned by the F-Learn is beneficial for the topology inference for roof reconstruction.

#### 4.3.2 F-Learn Strategy at Different Layers

We investigate the effect of applying the F-Learn strategy to feature maps of different levels with the HEAT model and the hierarchical ResNet50 backbone. The F-learn strategy is deployed to process feature maps generated at the first convolution layer, the first residual learning stage, and the second one. The sizes of the corresponding feature maps are  $128 \times 128$ ,  $64 \times 64$ , and  $32 \times 32$ , respectively. The quantitative results are presented in Tab. 4.

Table 4. **Quantitative comparison of applying the F-Learn strategy to features at different levels.** L0: the first convolution layer; L1: the first residual learning stage; L2: the second residual learning stage. Size: the feature map size. The F1-score metric is used for evaluating the performance of corner detection, edge inference, and region recovery. Higher scores mean better results. The **bold** values represent the best performance. (unit:%)

Method	Size	Corner	Edge	Region
HEAT	-	87.1	76.1	70.0
F-Learn-L0	$128 \times 128$	<b>88.6</b>	<b>79.1</b>	<b>73.4</b>
F-Learn-L1	$64 \times 64$	87.7	76.4	70.9
F-Learn-L2	$32 \times 32$	86.7	75.3	69.0

Tab. 4 implies that the higher the resolution is, the better the F-Learn strategy works. It is because high-resolution low-level feature maps contain more high-frequency geometric clues, while the low-resolution high-level ones are richer in abstract semantic information but poorer in geometric details. Therefore, the high-resolution feature maps are more suitable than the low-resolution counterparts for the F-Learn strategy to learn structurally informative geometries in roof reconstruction.

#### 4.3.3 F-Learn Strategy with Various Backbones

To further demonstrate the effectiveness of the F-learn strategy, we evaluate the performance improvement brought by the F-Learn strategy with different backbones. The quantitative results are presented in Tab. 5. As shown by Tab. 5, with the deployment of the F-learn strategy, the performance of HEAT is consistently improved in terms of corners detection, edge extraction, and region reconstruction.



Table 5. **Quantitative comparison of F-Learn strategy with HEAT built on different backbones, including ResNet18, ResNet34, and ResNet50.** Prec and F1 are the abbreviations of the precision and f1-score metrics. The higher the scores are, the better the performance is. The best results are marked **bold**. (unit:%)

	Backbone	Corner			Edge			Region		
		Prec	Recall	F1	Prec	Recall	F1	Prec	Recall	F1
w/o F-Learn	ResNet18	91.6	81.9	86.5	79.9	70.1	74.7	76.3	61.6	68.2
w/ F-Learn	ResNet18	<b>91.7</b>	<b>82.2</b>	<b>86.7</b>	<b>80.8</b>	<b>70.4</b>	<b>75.2</b>	<b>77.1</b>	<b>62.9</b>	<b>69.3</b>
w/o F-Learn	ResNet34	90.5	82.6	86.4	78.3	71.1	74.5	73.4	63.4	68.0
w/ F-Learn	ResNet34	<b>92.1</b>	<b>83.2</b>	<b>87.4</b>	<b>80.2</b>	<b>71.8</b>	<b>75.8</b>	<b>76.2</b>	<b>64.1</b>	<b>69.6</b>
w/o F-Learn	ResNet50	91.6	83.0	87.1	80.3	72.4	76.1	75.5	65.3	70.0
w/ F-Learn	ResNet50	<b>93.2</b>	<b>84.4</b>	<b>88.6</b>	<b>83.6</b>	<b>75.0</b>	<b>79.1</b>	<b>79.5</b>	<b>68.1</b>	<b>73.4</b>

#### 4.3.4 F-Learn Strategy in Topology Inference

In this section, we further explore the effectiveness of the F-Learn strategy in topology inference. Edges and regions are primitives that highly relate to the inference of topological relationships, and these primitives are learned with corners and low-level features in the HEAT model. Therefore, we directly use the ground-truth corner map to focus on topology inference. We compare our F-Learn strategy with the space-domain methods used in Sec. 4.3.1.

Table 6. **Quantitative results of the F-Learn and space-domain strategies with respect to topology inference.** The HEAT method with corner ground truth is set as the baseline. With corner annotations, all methods are studied with a focus on topology inference. The higher scores mean better results, and the best performance is **bolded**. (unit: %)

Evaluation Type →	Edge			Region		
	Prec	Recall	F1	Prec	Recall	F1
HEAT	93.5	84.1	88.5	90.2	73.8	81.2
HEAT-BConv	94.4	83.9	88.8	89.4	71.7	79.6
HEAT-Conv-Cas	94.8	85.3	89.8	<b>91.3</b>	73.7	81.6
HEAT-Conv-Par	93.8	84.8	89.1	89.7	73.9	81.0
<b>F-Learn</b>	<b>95.1</b>	<b>87.5</b>	<b>91.1</b>	89.5	<b>77.6</b>	<b>83.1</b>

Tab. 6 presents the numeric results, and it can be seen that our F-Learn strategy brings the highest gains in most metrics for edge inference and region reconstruction. This phenomenon demonstrates that the F-Learn strategy effectively supports topology inference with an outstanding ability for holistic geometry learning.

## 5. Conclusions

In this paper, we present a frequency-domain feature learning strategy (F-Learn) to tackle the issue of wrong topology recovery caused by the lack of holistic clues in low-level features. Experiments with a geometry recovery task convincingly validate the efficiency of the F-Learn

strategy in learning holistic geometry. In terms of the real scene, the F-Learn strategy achieves significant performance improvement of topological principle inference for roof structure reconstruction. The ablation studies verify that the F-Learn strategy outperforms the space-domain learning counterparts in capturing holistic geometric features regarding complex real scenes. We believe that it is promising to further explore the holistic learning ability brought by frequency analysis in more vision tasks.

## 6. Acknowledgement

This research is supported by NSFC-projects under Grant 42071370, the Fundamental Research Funds for the Central Universities of China under Grant 2042022dx0001, and the Open fund of Key Laboratory of Urban Land Resources Monitoring and Simulation, Ministry of Natural Resources under Grant KF202106084.

## References

- [1] Ronen Basri, Meirav Galun, Amnon Geifman, David Jacobs, Yoni Kasten, and Shira Kritchman. Frequency bias in neural networks for input of non-uniform density. In *International Conference on Machine Learning*, pages 685–694. PMLR, 2020. 3
- [2] Jiacheng Chen, Yiming Qian, and Yasutaka Furukawa. Heat: Holistic edge attention transformer for structured reconstruction. In *Proceedings of CVPR*, pages 3866–3875, 2022. 1, 3, 5, 6
- [3] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Q Weinberger, and Yixin Chen. Compressing convolutional neural networks in the frequency domain. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1475–1484, 2016. 3
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. 5
- [5] Ruoxi Deng, Chunhua Shen, Shengjun Liu, Huibing Wang, and Xinru Liu. Learning to predict crisp boundaries. In *ECCV*, pages 570–586, 2018. 2

- [6] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6569–6578, 2019. 2
- [7] Richard O Duda and Peter E Hart. Use of the hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1):11–15, 1972. 2
- [8] Tarik Dzanic, Karan Shah, and Freddie Witherden. Fourier spectrum discrepancies in deep network generated images. *Advances in neural information processing systems*, 33:3022–3032, 2020. 3
- [9] Rafael C Gonzales and Paul Wintz. *Digital image processing*. Addison-Wesley Longman Publishing Co., Inc., 1987. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the CVPR*, pages 770–778, 2016. 1, 5
- [11] Linxi Huan, Nan Xue, Xianwei Zheng, Wei He, Jianya Gong, and Gui-Song Xia. Unmixing convolutional features for crisp edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10):6602–6609, 2021. 2
- [12] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Fsd: Frequency space domain randomization for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6891–6902, 2021. 3
- [13] Jiaying Huang, Dayan Guan, Aoran Xiao, and Shijian Lu. Rda: Robust domain adaptation via fourier adversarial attacking. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8988–8999, 2021. 3
- [14] Kun Huang, Yifan Wang, Zihan Zhou, Tianjiao Ding, Shenghua Gao, and Yi Ma. Learning to parse wireframes in images of man-made environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 626–635, 2018. 2
- [15] Liming Jiang, Bo Dai, Wayne Wu, and Chen Change Loy. Focal frequency loss for image reconstruction and synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13919–13929, 2021. 3
- [16] Varsha Kamat-Sadekar and Subramaniam Ganesan. Complete description of multiple line segments using the hough transform. *Image and Vision Computing*, 16(9-10):597–613, 1998. 2
- [17] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European conference on computer vision (ECCV)*, pages 734–750, 2018. 2
- [18] Chen-Yu Lee, Vijay Badrinarayanan, Tomasz Malisiewicz, and Andrew Rabinovich. Roomnet: End-to-end room layout estimation. In *Proceedings of the ICCV*, pages 4865–4874, 2017. 1, 3
- [19] Hongjie Li, Xianwei Zheng, Mingyue Dong, Gui-Song Xia, and Hanjiang Xiong. Locally nonlinear affine verification for multisensor image matching. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–16, 2021. 1
- [20] Yancong Lin, Silvia L Pintea, and Jan C van Gemert. Deep hough-transform line priors. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, pages 323–340. Springer, 2020. 2
- [21] Yun Liu, Ming-Ming Cheng, Xiaowei Hu, Jia-Wang Bian, Le Zhang, Xiang Bai, and Jinhui Tang. Richer convolutional features for edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(8):1939–1946, 2019. 2
- [22] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2906–2917, 2021. 3
- [23] Charlie Nash, Yaroslav Ganin, SM Ali Eslami, and Peter Battaglia. Polygen: An autoregressive generative model of 3d meshes. In *ICML*, pages 7220–7229. PMLR, 2020. 1
- [24] Nelson Nauata and Yasutaka Furukawa. Vectorizing world buildings: Planar graph reconstruction by primitive detection and relationship inference. In *ECCV 2020*, pages 711–726. Springer, 2020. 1, 5, 6
- [25] Zequn Qin, Pengyi Zhang, Fei Wu, and Xi Li. Fcanet: Frequency channel attention networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 783–792, 2021. 3
- [26] Lawrence G Roberts. Machine perception of three-dimensional solids. *PhD thesis, Massachusetts Institute of Technology*, 1963. 2
- [27] Benoit Steiner, Zachary Devito, Soumith Chintala, Sam Gross, Adam Paszke, Francisco Massa, Adam Lerer, Gregory Chanan, Zeming Lin, Edward Yang, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, pages 8026–8037, 2019. 5
- [28] Sinisa Stekovic, Mahdi Rad, Friedrich Fraundorfer, and Vincent Lepetit. Montefloor: Extending mcts for reconstructing accurate large-scale floor plans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16034–16043, 2021. 3
- [29] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020. 3
- [30] Adam Van Etten, Dave Lindenbaum, and Todd M Bacastow. Spacenet: A remote sensing dataset and challenge series. *arXiv preprint arXiv:1807.01232*, 2018. 5
- [31] Yunhe Wang, Chang Xu, Chao Xu, and Dacheng Tao. Packing convolutional neural networks in the frequency domain. *IEEE transactions on pattern analysis and machine intelligence*, 41(10):2495–2510, 2018. 3
- [32] Rundi Wu, Chang Xiao, and Changxi Zheng. Deepcad: A deep generative network for computer-aided design models. In *Proceedings of the ICCV*, pages 6772–6782, 2021. 1, 3
- [33] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *Proceedings of the ECCV*, pages 466–481, 2018. 3
- [34] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. *International Journal of Computer Vision*, 125(1):3–18, 2017. 2

- [35] Kai Xu, Minghai Qin, Fei Sun, Yuhao Wang, Yen-Kuang Chen, and Fengbo Ren. Learning in the frequency domain. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1740–1749, 2020. 3
- [36] Yifan Xu, Weijian Xu, David Cheung, and Zhuowen Tu. Line segment detection using transformers without edges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4257–4266, 2021. 6
- [37] Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019. 3
- [38] 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
- [39] Nan Xue, Tianfu Wu, Song Bai, Fudong Wang, Gui-Song Xia, Liangpei Zhang, and Philip HS Torr. Holistically-attracted wireframe parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2788–2797, 2020. 6
- [40] Jimei Yang, Brian Price, Scott Cohen, Honglak Lee, and Ming-Hsuan Yang. Object contour detection with a fully convolutional encoder-decoder network. In *CVPR*, pages 193–202, 2016. 2
- [41] Zhiliang Zeng, Xianzhi Li, Ying Kin Yu, and Chi-Wing Fu. Deep floor plan recognition using a multi-task network with room-boundary-guided attention. In *Proceedings of the ICCV*, pages 9096–9104, 2019. 3
- [42] Fuyang Zhang, Nelson Nauata, and Yasutaka Furukawa. Conv-mpn: Convolutional message passing neural network for structured outdoor architecture reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2798–2807, 2020. 1, 3, 6
- [43] Fuyang Zhang, Xiang Xu, Nelson Nauata, and Yasutaka Furukawa. Structured outdoor architecture reconstruction by exploration and classification. In *Proceedings of ICCV*, pages 12427–12435, 2021. 3, 6
- [44] Xianwei Zheng, Zhuang Yuan, Zhen Dong, Mingyue Dong, Jianya Gong, and Hanjiang Xiong. Smoothly varying projective transformation for line segment matching. *ISPRS Journal of Photogrammetry and Remote Sensing*, 183:129–146, 2022. 1
- [45] Yichao Zhou, Haozhi Qi, and Yi Ma. End-to-end wireframe parsing. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 962–971, 2019. 2
- [46] Christian Zimmermann and Thomas Brox. Learning to estimate 3d hand pose from single rgb images. In *Proceedings of the ICCV*, pages 4903–4911, 2017. 3
- [47] Chuhan Zou, Alex Colburn, Qi Shan, and Derek Hoiem. Layoutnet: Reconstructing the 3d room layout from a single rgb image. In *Proceedings of the CVPR*, pages 2051–2059, 2018. 1, 3