

Supplementary Materials for Learning to Parse Wireframes in Images of Man-Made Environments

Kun Huang¹, Yifan Wang¹, Zihan Zhou², Tianjiao Ding¹, Shenghua Gao¹, and Yi Ma³

¹ShanghaiTech University {huangkun, wangyf, dingtj, gaoshh}@shanghaitech.edu.cn

²The Pennsylvania State University zzhou@ist.psu.edu

³University of California, Berkeley yima@eecs.berkeley.edu

1. Wireframe Construction Algorithm Detail

Given an image, our wireframe construction algorithm takes a set of junctions $\{\mathbf{p}_i\}_{i=1}^N$, $\mathbf{p}_i = (\mathbf{x}_i, \{\theta_{ik}\}_{k=1}^{K_i})$, and a line heat map h as input. Note that for the junctions and their branches predicted by our network, we only keep those with confidence scores higher than certain thresholds τ_c and τ_b , respectively. As a pre-processing step, we further adopt a strategy similar to non-maximum suppression to remove duplicate detections.

Our wireframe construction algorithm is presented in Alg. 1. In the algorithm, we first apply a threshold ω to convert the line heat map $h(p)$ into a binary map \mathcal{M} (line 2). Note that this threshold ω is varied to obtain the precision-recall curve in our experiments on wireframe construction. The algorithm then proceeds as follows:

First, we connect all pairs of junctions which are aligned with each other’s branch directions (lines 3-22). Let \mathbf{r}_{ik} represent the ray starting at \mathbf{p}_i along its k -th branch. We collect all possible rays as $\mathcal{R} = \{\mathbf{r}_{11}, \dots, \mathbf{r}_{1K_1}, \dots, \mathbf{r}_{i1}, \dots, \mathbf{r}_{iK_i}, \dots\}$, and use $(i, k) = \pi(t)$ to map the t -th ray in \mathcal{R} to its junction index i and branch index k . Then, for the rays in \mathcal{R} , we use $\mathcal{V} \in \mathbb{R}^{N_r \times N_r}$, $N_r = |\mathcal{R}|$, to record the indices of the corresponding ray/branch of the closest opposite junction. Specifically, $\forall t_1 \in \{1, \dots, N_r\}$, we set $\mathcal{V}(t_1, t_2)$ to 1 if and only if (i) \mathbf{p}_i is on the ray \mathbf{r}_{jk_2} and \mathbf{p}_j is on the ray \mathbf{r}_{ik_1} , where $(i, k_1) = \pi(t_1)$, $(j, k_2) = \pi(t_2)$, and (ii) the distance between \mathbf{p}_i and \mathbf{p}_j is the shortest among all such aligned pairs (lines 5-15). Then, we consider two rays are matched if $\mathcal{V}(t_1, t_2) = \mathcal{V}(t_2, t_1) = 1$ and add the corresponding junctions and line segments to \mathbf{P} and \mathbf{L} , respectively (lines 17-21).

Second, for any ray \mathbf{r}_{ik} which fails to find a matching ray using the above procedure, we attempt to recover additional line segments using the line support \mathcal{M} (lines 23-38). We consider the following cases:

- (a) If the distance between \mathbf{p}_i and \mathbf{q}_b , the intersection of \mathbf{r}_{ik} and the image boundary, is smaller than certain

threshold (say $0.05 \times m$ where m is the maximum of image width and height), we add $\{\mathbf{p}_i, \mathbf{q}_b\}$ and the connecting line segment to \mathbf{P} and \mathbf{L} , respectively (lines 24-26).

- (b) For a ray exceeding the length threshold in (a), we first find the farthest line pixel $\mathbf{q}_{\mathcal{M}}$ along the ray on \mathcal{M} . Then, we find all the intersection points $\{\mathbf{q}_1, \dots, \mathbf{q}_S\}$ of line segment $(\mathbf{p}_i, \mathbf{q}_{\mathcal{M}})$ with existing segments in \mathbf{L} (lines 28-29). Let $\mathbf{q}_0 = \mathbf{p}_i$ and $\mathbf{q}_{S+1} = \mathbf{q}_{\mathcal{M}}$, we calculate the line support ratio $\kappa(\mathbf{q}_{s-1}, \mathbf{q}_s)$, $s = \{1, \dots, S, S+1\}$, for each segment. Here, κ is defined as the ratio of line pixels (pixel p is a line pixel if $\mathcal{M}(p) = 1$) to the total length of the segment. If κ is above a threshold, say 0.6, we add the segment to \mathbf{L} and its endpoints to \mathbf{P} (lines 30-36).

2. Additional Experiments

2.1. Experiment on Junction Detection Network Parameters

In this section, we examine the choices of two important hyper-parameters in our junction detection network.

Effect of balancing positive and negative samples. In this experiment, we vary the value r_{\max} , which controls the maximum ratio between negative and positive samples at each iteration. Note that setting $r_{\max} = \infty$ is equivalent to using all grid cells during training. We can observe in Figure 1(a) that the precision-recall curves largely overlap. But as r_{\max} increases, the curve shifts toward the high-precision-low-recall regime, and vice versa. For example, when $r_{\max} = 1$, the precision and recall at $\tau = 0.5$ are 0.19 and 0.94, respectively. And when $r_{\max} = \infty$, the precision and recall at $\tau = 0.5$ are 0.70 and 0.44, respectively. Note that this has an important implication in practice, as human annotators tend to miss true junctions much more often than labelling wrong junctions. Empirically, we have found that $r_{\max} = 7$ yields more satisfactory results.

Algorithm 1 Wireframe Construction

Input: Junctions $\{p_i\}_{i=1}^N$, $p_i = (\mathbf{x}_i, \{\theta_{ik}\}_{k=1}^{K_i})$, and a line heat map $h(p)$

Output: Wireframe \mathcal{W} consisting of a set of junction points \mathbf{P} connected by a set of line segments \mathbf{L}

```
1: Initialize  $\mathbf{P} \leftarrow \emptyset$ ,  $\mathbf{L} \leftarrow \emptyset$ ,  $\mathcal{V} \leftarrow \mathbf{0}$ 
2: Binarize  $h(p)$  with threshold  $\omega$  into  $\mathcal{M}$ 
3: for  $t_1 \in \{1, 2, \dots, N_r\}$  do
4:    $(i, k_1) \leftarrow \pi(t_1)$ ,  $d_{\min} \leftarrow \infty$ ,  $z \leftarrow 0$ 
5:   for  $t_2 \in \{1, 2, \dots, N_r\}$  do
6:      $(j, k_2) \leftarrow \pi(t_2)$ 
7:     if  $j \neq i$  and  $p_j$  on  $r_{ik_1}$  and  $p_i$  on  $r_{jk_2}$  then
8:       if  $\|\mathbf{x}_i - \mathbf{x}_j\| < d_{\min}$  then
9:          $d_{\min} \leftarrow \|\mathbf{x}_i - \mathbf{x}_j\|$ ,  $z \leftarrow t_2$ 
10:      end if
11:    end if
12:  end for
13:  if  $z \neq 0$  then
14:     $\mathcal{V}(t_1, z) \leftarrow 1$ 
15:  end if
16: end for
17: for all  $t_1, t_2 \in \{1, 2, \dots, N_r\}$ ,  $t_1 \neq t_2$  do
18:   if  $\mathcal{V}(t_1, t_2) = 1$  and  $\mathcal{V}(t_2, t_1) = 1$  then
19:      $(i, k_1) \leftarrow \pi(t_1)$ ,  $(j, k_2) \leftarrow \pi(t_2)$ 
20:      $\mathbf{P} \leftarrow \mathbf{P} \cup \{p_i, p_j\}$ ,  $\mathbf{L} \leftarrow \mathbf{L} \cup \{(p_i, p_j)\}$ 
21:   end if
22: end for
23: for all  $r_{ik}$  not matched to another ray do
24:   Find the intersection of  $r_{ik}$  and image boundary  $q_b$ 
25:   if  $\|\mathbf{x}_i - q_b\| \leq 0.05 \times m$  then
26:      $\mathbf{P} \leftarrow \mathbf{P} \cup \{p_i, q_b\}$ ,  $\mathbf{L} \leftarrow \mathbf{L} \cup \{(p_i, q_b)\}$ 
27:   else
28:     Find the farthest point  $q_{\mathcal{M}}$  along  $r_{ik}$  on  $\mathcal{M}$ 
29:     Find all intersections  $\{q_1, \dots, q_S\}$  of  $(p_i, q_{\mathcal{M}})$ 
    with segments in  $\mathbf{L}$ 
30:      $q_0 \leftarrow p_i$ ,  $q_{S+1} \leftarrow q_{\mathcal{M}}$ 
31:     for  $s \in \{1, 2, \dots, S, S+1\}$  do
32:       if  $\kappa(q_{s-1}, q_s) > 0.6$  then
33:          $\mathbf{P} \leftarrow \mathbf{P} \cup \{q_{s-1}, q_s\}$ 
34:          $\mathbf{L} \leftarrow \mathbf{L} \cup \{(q_{s-1}, q_s)\}$ 
35:       end if
36:     end for
37:   end if
38: end for
```

Going deeper. It is also interesting to investigate how the network depth of the encoder affects the performance. In this experiment, we compared two different choices based on Google Inception-v2, namely the first layer to “Mixed_3b”, and the first layer to “Mixed_4b”. Note that the latter has a larger depth and receptive field, at the cost of spatial resolution (30×30). As one can see in Figure 1(b),

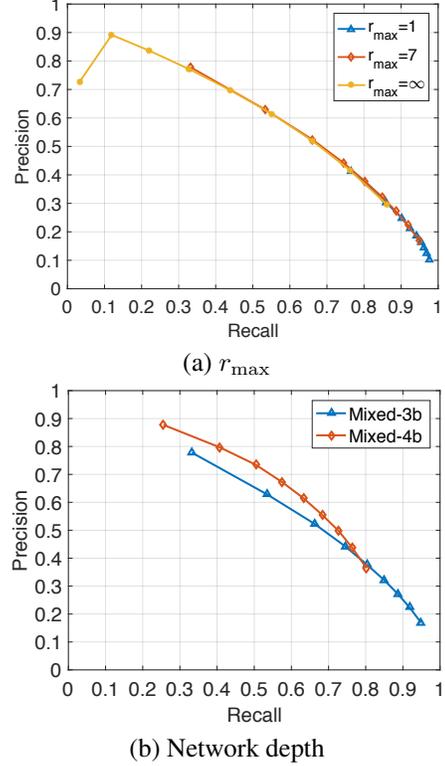


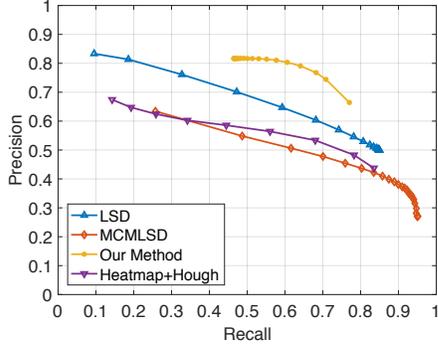
Figure 1. Experiment on junction detection network parameters.

increasing the depth (i.e., predicting at the “coarser” level) results in higher precision but lower recall. This suggests possibilities to further improve the performance of our method using a “skip-net” architecture, that is, combining predictions at multiple levels. We leave this for future work.

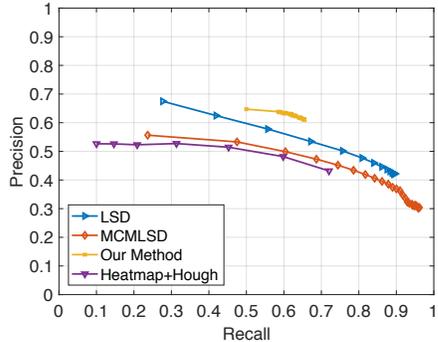
2.2. Experiment on Line Segment Detection

In this experiment, we study the possibility of extracting line segments *directly* from the pixel-wise line heat map predicted by our network (i.e., without using junctions). To this end, we simply perform a probabilistic hough transform [1] on the line heat map to generate line segments. We compare the results with LSD, MCMLSD, and our full wireframe construction method.

Figure 2 shows the precision-recall curves of all methods. We make the following observations on the results: *First*, the performance of our “Heatmap + Hough” approach is comparable to that of the state-of-the-art line segment detection method MCMLSD, verifying the effectiveness of our line detection network. *Second*, by combining the predicted junctions with the line heat map, our full wireframe construction method performs significantly better than using the line heat map alone. This further illustrates the importance of junction detection in parsing the wireframe: By detecting the “endpoints” of the line segments, we effectively overcome the difficulties faced by tra-



(a) Our test dataset



(b) York Urban dataset

Figure 2. Experiment on line segment detection.

ditional line segment detection methods, including the false detection problem and the inaccurate endpoint problem.

2.3. Additional Results on Junction Detection

In Figure 4, we show additional junction detection results obtained by all methods. One can see that our method is able to detect most junctions and their branches in the image, achieving superior performance over existing methods.

From Figure 4 we can also observe some limitations of our method. Specifically, there are occasionally repeated detections in our result. This may be caused by junctions located at the boundary of two adjacent grid cells used in our junction detection network. Similarly, the use of grid could also lead to missed detection if two junctions are very close to each other. But we note that such cases are rather uncommon in practice and have very small effect on the overall scene structure estimation.

2.4. Additional Results on Wireframe Construction

In Figure 5, we show additional wireframe detection results obtained by all methods. Our method outperforms other two in most areas and produces much cleaner results as we focus on long line segments and exploit their relations (junctions). Therefore, the resulted wireframes are potentially more suitable for 3D reconstruction tasks.

In Figure 3, we further show some failure cases of our

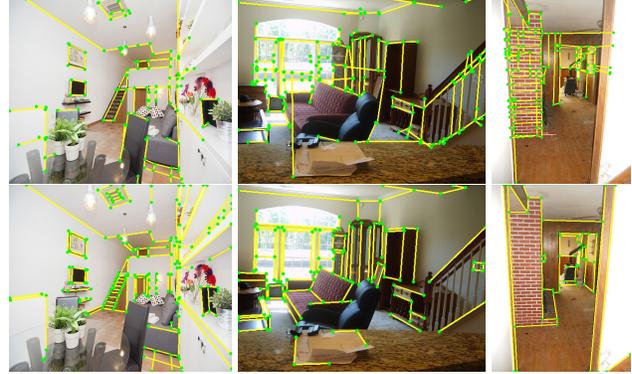


Figure 3. Failure cases on our test dataset. **First row:** Our method. **Second row:** Ground truth.

method. One challenging case corresponds to structures with relatively small scale and weak image gradients (e.g., the stairs in the first image). Also, our method sometimes has difficulty in image region of repetitive patterns (e.g., the handrails in the second image and the brick wall in the third image), generating fragment, incomplete results. This suggests opportunities for further improvement by explicitly harnessing such geometric structure in our wireframe construction.

References

- [1] J. Matas, C. Galambos, and J. Kittler. Robust detection of lines using the progressive probabilistic hough transform. *Computer Vision and Image Understanding*, 78(1):119–137, 2000. 2



Figure 4. Junction detection results. **First row:** MJ ($d_{\max} = 20$). **Second row:** ACJ ($\epsilon = 1$). **Third row:** Our method ($\tau = 0.5$).

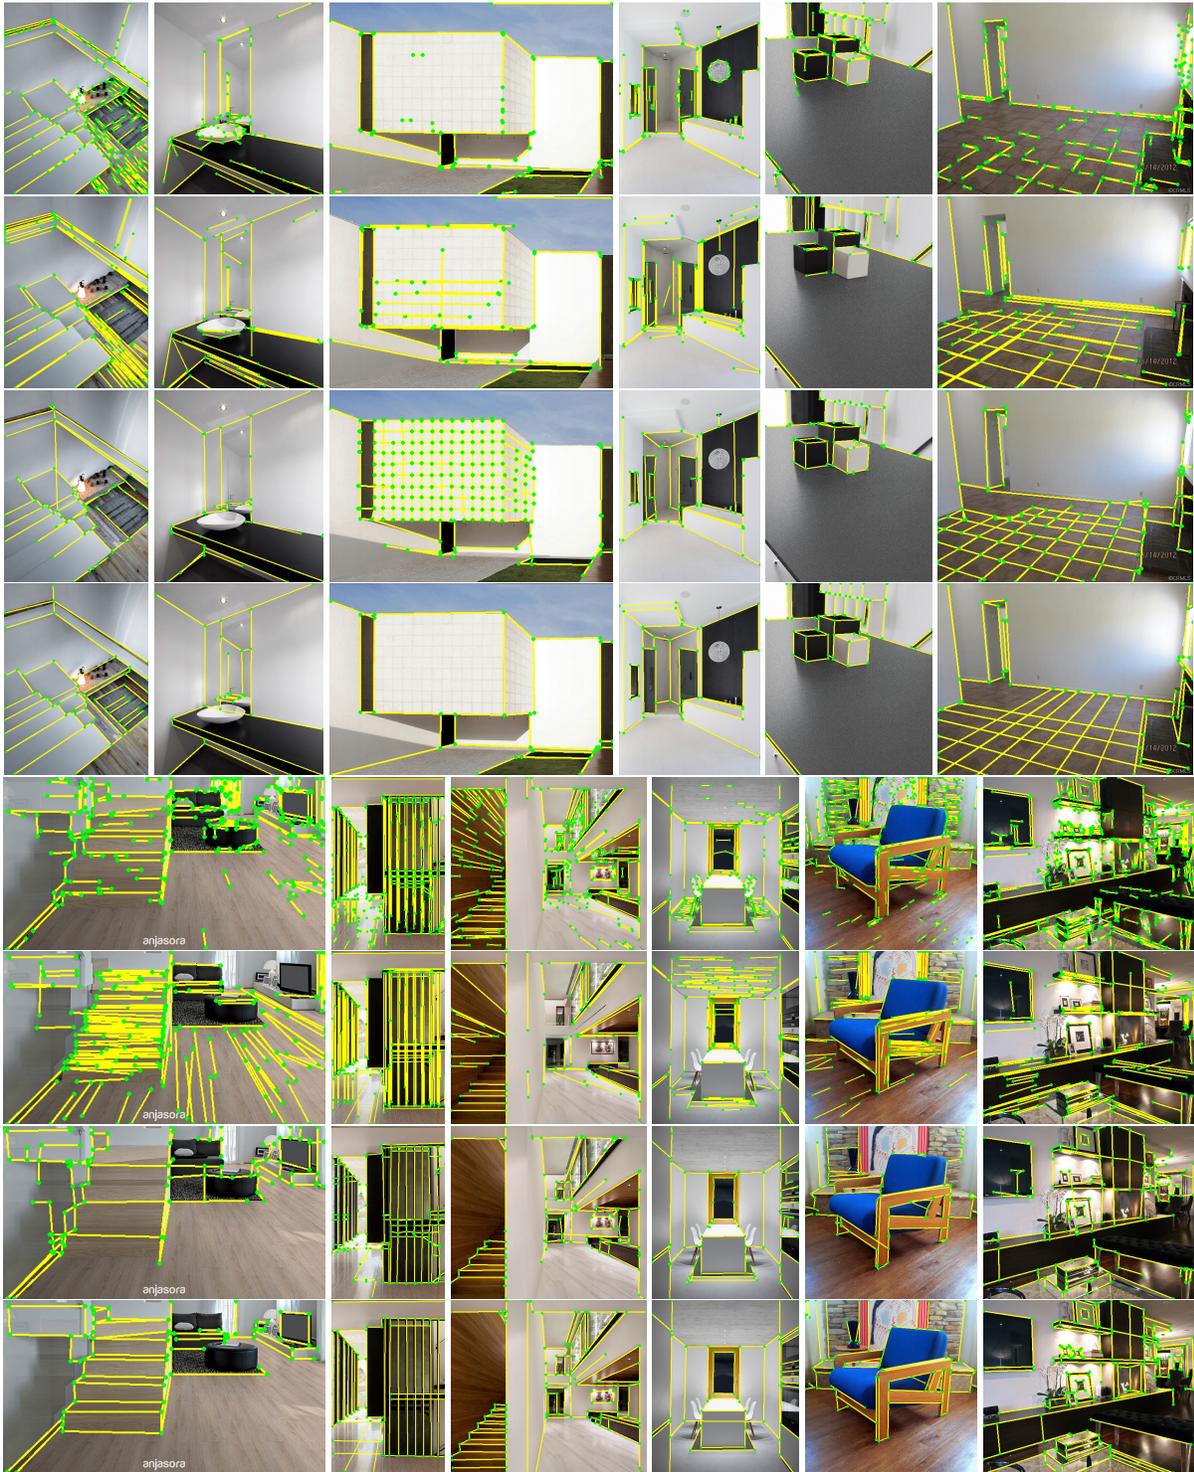


Figure 5. Line/wireframe detection results. **First row:** LSD ($-\log(\text{NFA}) > 0.01 \times 1.75^8$). **Second row:** MCMLSD (confidence top 100). **Third row:** Our method (line heat map $h(p) > 10$). **Fourth row:** Ground truth.