

# Improving Transformer Based Line Segment Detection with Matched Predicting and Re-ranking

Xin Tong, Shi Peng, Baojie Tian, Yufei Guo, Xuhui Huang, Zhe Ma\*

Intelligent Science & Technology Academy of CASIC  
xin\_tong@pku.edu.cn, mazhe.thu@163.com

## Abstract

Classical Transformer-based line segment detection methods have delivered impressive results. However, we observe that some accurately detected line segments are assigned low confidence scores during prediction, causing them to be ranked lower and potentially suppressed. Additionally, these models often require prolonged training periods to achieve strong performance, largely due to the necessity of bipartite matching. In this paper, we introduce RANK-LETR, a novel Transformer-based line segment detection method. Our approach leverages learnable geometric information to refine the ranking of predicted line segments by enhancing the confidence scores of high-quality predictions in a posterior verification step. We also propose a new line segment proposal method, wherein the feature point nearest to the centroid of the line segment directly predicts the location, significantly improving training efficiency and stability. Moreover, we introduce a line segment ranking loss to stabilize rankings during training, thereby enhancing the generalization capability of the model. Experimental results demonstrate that our method outperforms other Transformer-based and CNN-based approaches in prediction accuracy while requiring fewer training epochs than previous Transformer-based models.

## Introduction

Line segment detection is a fundamental and critical problem in computer vision. An accurate line segment detection algorithm can significantly enable and enhance various computer vision applications, such as 3D reconstruction (Li, Liu, and Zhou 2012; Langlois, Boulch, and Marlet 2019), camera calibration (Zhang et al. 2016; Nakano 2021), depth estimation (Zavala and Martinez-Carranza 2022), scene understanding (Hofer, Maurer, and Bischof 2017), object detection (Tang et al. 2022), and SLAM (Vakhitov and Lempitsky 2019; Gomez-Ojeda et al. 2019). Traditional line segment detection algorithms directly utilize low-level information, such as image gradients, for line segment detection. While these algorithms offer fast detection speeds, they often result in fragmented line segments. In contrast, learning-based methods can detect longer and more meaningful line segments.

\*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Comparison between unmatched and matched predicting in Transformer-based line segment detection. (a) In unmatched predicting, predictions need to be matched with ground truth online using bipartite matching. (b) In matched predicting, each prediction is directly associated with a specific ground truth value, resulting in higher training efficiency and stability.

Convolutional Neural Network (CNN)-based methods typically predict the centroid of a line segment along with the offsets of its two endpoints relative to the centroid, taking advantage of the translation invariance of convolution operations. In contrast, Transformer-based methods can directly predict the two endpoints of a line segment. However, existing Transformer-based line segment detection methods employ Transformer encoder-decoder architecture, which requires online bipartite matching to define the loss function during training. As a result, these methods often demand extensive training time to converge and depend on heavy pre-training or multi-stage fine-tuning to achieve optimal results. To address this issue, we propose a novel matched prediction strategy, where the matching correspondence between predictions and ground truths is predefined, as shown in Fig.1. Specifically, the feature point closest to the centroid of a line segment is responsible for predicting that line. The predictions are then ordered and can directly correspond to the ground truth supervision signal.

Proposal-based methods typically assign a score to each predicted line segment and rank them from highest to lowest, with the top-ranked segments chosen as the final predictions. In these methods, scores and locations are predicted simultaneously, making score prediction and location regres-

sion independent processes. However, we observe that some accurately detected line segments receive low confidence scores during prediction, leading to lower rankings and potential suppression. This observation suggests the possibility of reevaluating and optimizing the scores of predicted line segments to improve the final selection. As shown in Table 1, our experiments indicate that selecting the appropriate line segments from the candidates can significantly enhance detection performance. To address this, we present a simple and efficient re-ranking module for line segment detection based on learnable geometric information, since low-level geometric info such as edge, endpoint, and length of line segment always provides essential insights in finding better lines. Moreover, the re-ranking module is highly interpretable and brings minimal computational overhead.

In this paper, we propose a novel Transformer-based method for line segment detection. The multi-scale and multi-level image features from different layers of the CNN backbone with feature pyramid are extracted, and then processed by deformable Transformer Encoder. The scores and locations of line segments are predicted on higher-resolution feature maps. Specifically, each feature point predicts a confidence score indicating the likelihood that the centroid of a given line segment is nearest to that feature point, as well as the location of the potential line segment. To supervise the training, we directly apply confidence and position losses without the need for bipartite matching. Additionally, we introduce a novel ranking loss to ensure that feature points predicting higher-quality line segments receive higher confidence scores. Finally, the line segments are re-ranked by optimizing the scores based on learnable geometric information, with the top-ranked segments selected as the final predictions.

Our contributions can be summarized as follows:

(1) We propose a novel Transformer-based line segment detection architecture where the feature point nearest to the centroid of a line segment is responsible for predicting its location. This architecture eliminates the need for bipartite matching and enables line segment detection on high-resolution feature maps.

(2) We observe the confidence scores may fail to accurately reflect the quality of the predicted line segments. To tackle the problem, we introduce a learnable geometric information based re-ranking module and a line segment ranking loss, which can significantly improve the detection performance.

(3) Experimental results demonstrate that our method outperforms other Transformer-based and CNN-based approaches in prediction accuracy while requiring fewer training epochs than previous Transformer-based models.

## Related Works

**Line Segment Detection.** Traditional line detection methods such as (Von Gioi et al. 2008; Akinlar and Topal 2011; Lu et al. 2015) rely on low-level image features, e.g., image gradients. Based on local edge features, Hough transform is used for line segment detection in (Guil, Villalba, and Zapata 1995; Furukawa and Shinagawa 2003; Xu, Shin, and Klette 2015). Recently, learning-based methods have

Predictions	sAP <sup>5</sup> (%)	sAP <sup>10</sup> (%)
w/o Re-ranking	62.0	67.2
Re-ranking with GT	91.3 (+ 29.3)	95.5 (+ 28.3)

Table 1: Line detection performance before and after re-ranking using ground truth. Predicted line segments that are closer to the ground truth are assigned higher scores through re-ranking. With proper ranking, the overall performance of the predicted candidate line segments can be significantly improved. Motivated by this observation, we propose to use a re-ranking module to optimize the predicted confidence scores during inferring and use a ranking loss to supervise confidence score predicting additionally during training.

achieved promising results and can be roughly divided into two categories. In junction-based methods, DWP (Huang et al. 2018) predicts junction map and edge map in two branches before merging them. PPGNet (Zhang et al. 2019) uses a point-pair graph to describe junctions and line segments. L-CNN (Zhou, Qi, and Ma 2019) applies line proposal and LoI pooling to propose candidate lines and verify them. LETR (Xu et al. 2021) models it as object detection and predicts line segments with DETR architecture. Methods with dense prediction first predict representation map and extract line segments with post-processing. AFM (Xue et al. 2019) proposes attraction field maps to represent the image space and uses a squeeze module to generate line segment maps. HAWP (Xue et al. 2020) further builds a hybrid model considering AFM and L-CNN. Lin *et al.* (Lin, Pintea, and van Gemert 2020) apply deep Hough transform to the previous detection architectures. TP-LSD (Huang et al. 2020b) introduces tri-points line segment representation for end-to-end detection. M-LSD (Gu et al. 2022) presents SoL augmentation and designs an extremely efficient architecture for fast detection. In this work, we propose a novel Transformer-based line detection method which can get proposals from dense feature maps without bipartite matching.

**Visual Transformer.** Transformer-based models have gained significant success in computer vision tasks. Dosovitskiy *et al.* (Dosovitskiy et al. 2020) propose the Vision Transformer (ViT) which applies transformers to image classification by dividing the image into fixed-size patches. Carion *et al.* (Carion et al. 2020) introduce transformers into object detection pipeline through bipartite matching, named DETR. Zhu *et al.* (Zhu et al. 2020) further propose deformable DETR whose attention modules only attend to a small set of keys. Xu *et al.* (Xu et al. 2021) apply DETR architecture in line segment detection with a multi-scale encoder-decoder strategy. Tong *et al.* (Tong et al. 2022) use Transformer to cluster line segments corresponding with the same vanishing points and further apply Transformer in end-to-end vanishing point detection (Tong et al. 2024). Transformers are also used in semantic segmentation (Xie et al. 2021), pose estimation (Huang et al. 2020a), tracking (Chen et al. 2021), etc. Our work aims at improving transformer-based line segment detection with matched predicting strategy and presented re-ranking module.

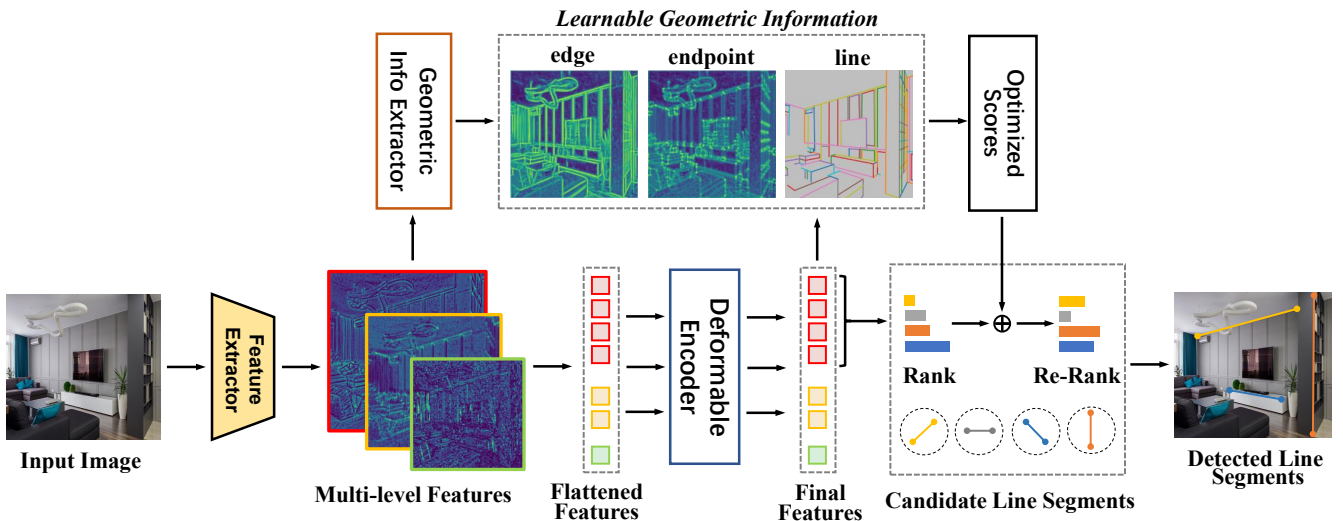


Figure 2: Overview of the proposed RANK-LETR. The process begins by feeding an image into a CNN backbone to extract multi-level feature maps from different layers. These features are then processed by a deformable Transformer encoder to generate candidate line segments. The candidate segments are predicted using high-resolution feature maps for higher prediction accuracy and less ambiguity, with each segment represented by confidence scores and positions. Each feature point is responsible for detecting the line segment whose centroid is nearest to it. Additionally, learnable geometric information is extracted from the multi-level features using a CNN-based geometric information extractor. Finally, the line segments are re-ranked by optimizing their confidence scores with the learnable geometric information.

**Ranking-based Losses for Object Detectors.** Ranking-based losses have been widely used recently. Average Precision Loss is first proposed (Chen et al. 2019) to address the imbalance of foreground-background classification problem by framing object detection as a ranking task. Oksuz *et al.* (Oksuz et al. 2021) propose Rank & Sort (RS) Loss that defines a ranking objective between positives and negatives as well as a sorting objective to prioritize positives with respect to their continuous IoUs. Yavuz *et al.* (Yavuz et al. 2024) apply Bucketed Ranking-based (BR) Losses which group negative predictions into several buckets. Cetinkaya *et al.* (Cetinkaya, Kalkan, and Akbas 2024) extend Ranking-based Loss to edge detection. In this work, we propose a line segment ranking loss for ranking the feature point with a higher line segment detection quality a higher confidence score.

## Method

The overview of our algorithm is depicted in Fig. 2. We introduce our method from three aspects including problem modeling, network architecture and training supervision.

### Line Segment Proposal for Matched Predicting

In traditional DETR frameworks, bipartite matching is used to assign predictions to ground truth sequences. However, these approaches can be unstable in convergence, particularly during the early stages of training (Li et al. 2022). For the task of line segment detection, we observe that the centroids of individual line segments often do not coincide. Based on this observation, we introduce a novel centroid-based representation for line segments in Transformer-based

predictions. Specifically, in proposal-based line segment detection, each feature point in the feature map has a unique 2D coordinate, and the feature point closest to the centroid of a line segment is responsible for predicting it. Leveraging the Transformer architecture, which incorporates positional information through positional encoding, each feature point can directly predict the endpoint coordinates  $(x_s, y_s, x_e, y_e)$  of the line segment, rather than predicting the offsets of the endpoints from the centroid as in CNN-based methods (Huang et al. 2020b; Gu et al. 2022).

Given the processed feature point  $\mathbf{p}$  in the feature map  $F$ , we can get the confidence score  $c$  indicating the confidence that the feature point should predict a line segment

$$c = \sigma(\mathcal{F}_c(\mathbf{p})), \quad (1)$$

and the location of the corresponding line segment

$$\mathbf{l} = \mathcal{F}_l(\mathbf{p}), \quad (2)$$

where  $\mathcal{F}_c, \mathcal{F}_l$  are CNN-based prediction modules and  $\sigma$  is the sigmoid function.

### Line Re-ranking Module

While confidence scores are assigned to predicted line segments, these scores may not fully reflect the true quality of the segments since they are generated simultaneously. We propose optimizing these scores by evaluating the predicted line segments against corresponding geometric information, such as edges, endpoints, and length. To accomplish this, we perform sampling-based evaluations separately for endpoints and edges using the generated line segments. Specifically, we create endpoint maps and edge maps with a geo-

metric information extractor. For a line segment to be considered high-quality, its endpoints should align with high-confidence regions on the endpoint map, and the segment should significantly overlap with high-confidence areas on the edge map.

Given the endpoint map  $M_e$ , we sample the endpoints  $e_1, e_2$  according to the predicted line segments. The sampled scores are averaged as the endpoint confidence, which can be represented as

$$s_e = \frac{1}{2} \sum \mathcal{S}(\{e_1, e_2\}; M_e), \quad (3)$$

$\mathcal{S}(\{e\}; M)$  represents the sampling kernel indicating sampling  $M$  with the location  $e$ . Given the edge map  $M_d$ , we first uniformly sample  $m$  points between the two endpoints of the line segments. The scores are then sampled with these points from the edge map, and finally averaged as the edge score, which can be represented as

$$s_d = \frac{1}{m} \sum \mathcal{S}(\{\frac{k}{m}e_1 + \frac{m-k}{m}e_2\}_{k \in [1, m]}; M_d). \quad (4)$$

The geometric information extractor is composed of a group of CNNs in this work, which is supervised by ground truth edge maps and endpoint maps generated from the ground truth line segments.

Moreover, to encourage the detection of long line segments and suppress the detection of fragmented line segments, the length is taken into account in evaluating line segment quality. To avoid the score expansion caused by length, the length scores is defined as

$$s_l = \ln(\|e_1 - e_2\|_2 + 1). \quad (5)$$

The optimized score is finally defined as the sum of the above terms and the previous confidence score  $c$ , which can be represented as

$$s = \delta_e s_e + \delta_d s_d + \delta_l s_l + c. \quad (6)$$

By incorporating the re-ranking method, we can refine the confidence scores of predicted line segments by considering not only the centroid confidence score but also the overall quality of the line segments, as indicated by their endpoints, overlap with the edge map and the length.

## Network Architecture

Our network architecture is built on a CNN-based backbone combined with a Deformable Transformer, designed specifically for line segment detection, which requires abundant low-level, high-resolution information. To capture features at various resolutions, we employ a feature pyramid during CNN feature extraction, feeding these multi-scale features into the Transformer. Notably, our approach simplifies the Transformer by using only the Encoder befitting from the stable line segment representation. Each feature token in the Encoder layers corresponds directly to a prediction target, eliminating the need for learnable queries and bipartite matching. For line segment prediction, we utilize the final  $128 \times 128$  high-resolution features, which allow for finer predictions and reduce conflicts from overlapping centroids.

Additionally, we enhance the network’s sensitivity to line segments by incorporating rotation augmentation during image feature extraction, rotating the images by  $\pm 90^\circ$  and applying an inverse rotation afterward. Our overall architecture includes an image feature extractor based on ResNet50 with a feature pyramid and a feature processing module consisting of 6 deformable Transformer Encoder layers.

## Training Supervision

Since our method maps predicted results to ground truth before prediction, bipartite matching is no longer required to establish correspondences. Instead, we directly employ confidence loss and position loss to supervise our line segment detection model during training. Specifically, for confidence loss, we use binary cross-entropy loss to determine whether a feature point should predict a line segment, i.e., whether the centroid of the line segment is closest to the feature point. We find that the number of feature points that do not need to be responsible for predicting the line segments (defined as negative feature points) is considerably larger than those that need to be responsible (defined as positive feature points). Thus, we randomly select some of the former for training. The confidence loss  $\mathcal{L}_{conf}$  is defined as

$$\mathcal{L}_{conf} = -\frac{1}{N_{\pm}} \sum_{i=1}^{N_{\pm}} (\hat{c}_i \log c_i + (1 - \hat{c}_i) \log(1 - c_i)), \quad (7)$$

where  $N_{\pm}$  is the total number of the positive and selected negative feature points.  $c_i$  is the predicted confidence score of the  $i$ -th feature point.  $\hat{c}_i$  is the ground truth indicates that whether the  $i$ -th feature point should respond to detect a line segment and is annotated as +1 or 0.

In the position loss, we minimize the L1 distance between the predicted line segment and ground truth ones on the image space. The loss is only applied to the positive feature points since the prediction of negative anchors should be meaningless. The presented position loss  $\mathcal{L}_{pos}$  can be represented as

$$\mathcal{L}_{pos} = \sum_{i=1}^{N_+} \|l_i - \hat{l}_i\|_1, \quad (8)$$

where  $N_+$  is the number of positive feature points.  $l_i$  represents a predicted line segments from the positive feature point and  $\hat{l}_i$  is the corresponding ground-truth.

Inspired by the success of recent ranking-based losses in detection tasks, we introduce a line segment ranking loss to ensure that feature points with higher line segment detection quality receive higher confidence scores. Many ranking-based losses are guided by the quality of the bounding box or the predicting confidence. In the context of line segment detection, the natural basis for ranking is the quality of the predicted line segment. We define this quality using the Euclidean distance between the predicted line segment and its corresponding ground truth. A shorter distance indicates higher prediction quality. For all positive feature points, we perform pairwise comparisons based on their prediction quality, ensuring that feature points with higher prediction quality receive higher confidence scores, which can

Method	Wireframe					YUD					Training Epochs	FPS
	sAP <sup>5</sup>	sAP <sup>10</sup>	sF <sup>10</sup>	sF <sup>15</sup>	LAP	sAP <sup>5</sup>	sAP <sup>10</sup>	sF <sup>10</sup>	sF <sup>15</sup>	LAP		
LSD (Von Gioi et al. 2008)	6.7	8.8	-	-	18.7	7.5	9.2	-	-	16.1	-	100.0
DWP (Huang et al. 2018)	3.7	5.1	-	-	6.6	2.8	2.6	-	-	3.1	120	2.2
AFM (Xue et al. 2019)	18.3	23.9	-	-	36.7	7.0	9.1	-	-	17.5	200	14.1
LGNN (Meng et al. 2020)	-	62.3	-	-	-	-	-	-	-	-	-	15.8
TP-LSD (Huang et al. 2020b)	57.6	57.2	-	-	61.3	<b>27.6</b>	27.7	-	-	<b>34.3</b>	400	20.0
LETR (Xu et al. 2021)	59.2	65.6	<u>66.1</u>	<u>67.4</u>	<u>65.1</u>	24.0	27.6	<u>39.6</u>	<b>41.1</b>	32.5	825	5.4
L-CNN (Zhou, Qi, and Ma 2019)	58.9	62.8	61.3	62.4	59.8	25.9	28.2	36.9	37.8	32.0	16	16.6
HAWP (Xue et al. 2020)	62.5	66.5	64.9	65.9	62.9	26.1	<u>28.5</u>	<b>39.7</b>	40.5	30.4	30	32.9
M-LSD (Gu et al. 2022)	56.4	62.1	-	-	61.5	24.6	27.3	-	-	30.7	150	115.4
M-LSD†(Gu et al. 2022)	<u>63.3</u>	<u>67.1</u>	-	-	64.2	<u>27.5</u>	<u>28.5</u>	-	-	32.4	150	32.9
RANK-LETR (Ours)	<b>65.0</b>	<b>69.7</b>	<b>66.7</b>	<b>67.7</b>	<b>65.6</b>	<b>27.6</b>	<b>30.1</b>	<b>39.7</b>	<u>40.6</u>	<u>34.1</u>	120	9

Table 2: Comparison results on Wireframe (Huang et al. 2018) and YUD (Denis, Elder, and Estrada 2008) datasets. We compare the proposed method with LSD (Von Gioi et al. 2008), DWP (Huang et al. 2018), AFM (Xue et al. 2019), LGNN (Meng et al. 2020), TP-LSD (Huang et al. 2020b), LETR (Xu et al. 2021), L-CNN (Zhou, Qi, and Ma 2019), HAWP (Xue et al. 2020) and M-LSD (Gu et al. 2022). M-LSD† denotes the method of combining M-LSD and HAWP. Our method outperforms existing methods in prediction accuracy while requiring fewer training epochs than previous Transformer-based models.

be represented as

$$\mathcal{L}_{rank} = -\frac{1}{N_+^2} \sum_{i=1}^{N_+} \sum_{j=1}^{N_+} \sigma((c_j - c_i)) \quad (9)$$

$$(\|\mathbf{l}_i - \hat{\mathbf{l}}_i\|_2 - \|\mathbf{l}_j - \hat{\mathbf{l}}_j\|_2),$$

where we use sigmoid function to approximate the step function.

Moreover, as we need to use the junction map and edge map in our line re-ranking module, we predict junction maps and edge maps from feature maps in different resolutions. For each resolution, we supervise the junction map  $\mathcal{F}_j(\mathbf{F})$  and edge map  $\mathcal{F}_e(\mathbf{F})$  as similar to ground truth maps  $J$  and  $E$ , which can be represented as

$$\mathcal{L}_{junc} = \sum_{i=1}^n \|\mathcal{F}_j(\mathbf{F}_i) - J_i\|_2, \quad (10)$$

and

$$\mathcal{L}_{edge} = \sum_{i=1}^n \|\mathcal{F}_e(\mathbf{F}_i) - E_i\|_2. \quad (11)$$

The total loss can be defined as the sum of the above loss terms with proper weights, which can be written as

$$\mathcal{L}_{total} = \lambda_r \mathcal{L}_{rank} + \lambda_c \mathcal{L}_{conf} + \lambda_p \mathcal{L}_{pos} + \lambda_j \mathcal{L}_{junc} + \lambda_e \mathcal{L}_{edge}. \quad (12)$$

## Experimental Results

In this section, we first describe the implementation setup of our method. Then we compare it with the state-of-the-art line segment detection approaches with quantitative experiments. In the ablation study, we demonstrate the effectiveness of the components in our method. We also conduct a parameter study to select better hyperparameters.

## Experimental Setup

**Datasets and Metrics** We conduct our experiments in two publicly available datasets including the Wireframe dataset (Huang et al. 2018) and the YorkUrban dataset (Denis, Elder, and Estrada 2008). The Wireframe dataset consists of 5,000 training and 462 test images of man-made environments, while the YorkUrban dataset has 102 test images. Following the typical training and test protocol (Huang et al. 2020b; Zhou, Qi, and Ma 2019), we train our model with the training set from the Wireframe dataset and test with both Wireframe and YorkUrban datasets. We evaluate our models using prevalent metrics for line segment detection task including structural average precision (sAP), the F-score measurement (sF) and line matching average precision (LAP).

**Implementation details** Our training and evaluation are implemented in PyTorch. We use 4 NVIDIA Tesla V100 cards for training and 1 card for evaluation. In training, We use AdamW as the model optimizer and set weight decay as  $10^{-4}$ . We train the model for 120 epochs. The initial learning rates are set to  $5 \times 10^{-4}$  for all parameters. Learning rates are reduced by a factor of 10 in epoch 60 and 90. We use a batch size of 8 and the size of the input images is set to  $512 \times 512$ . Similarly to many computer vision tasks, we adopt a CNN backbone pretrained on ImageNet, while other parameters are trained from scratch.

The results of our method is predicted on the features of the resolution of  $128 \times 128$ .  $\lambda_r, \lambda_c, \lambda_p, \lambda_j, \lambda_e$  are set to 1, 1, 10, 1, 1, respectively.  $m$  is set as 32 for sampling and  $\delta_e, \delta_d, \delta_l$  are set to 0.5, 0.5, 0.5 empirically for YUD dataset and 0.4, 0.1, 0.2 after validating on random sampling set on training set. 500 line segments with high scores are detected with NMS for the proposed method in all the experiences.

## Comparison with the SOTA

Our comparisons are conducted on the Wireframe dataset (Huang et al. 2018) and the YorkUrban dataset (Denis, Elder,



Figure 3: Visual examples of line segment detection results of two Transformer-based methods including LETR and ours on the Wireframe dataset. Our method can produce more accurate and complete detection results. For a better visual experience, we emphasize some examples of accurate detection with red bounding boxes and complete detection with purple bounding boxes.

and Estrada 2008). We compare our method with the state-of-the-art methods including LSD (Von Gioi et al. 2008), DWP (Huang et al. 2018), AFM (Xue et al. 2019), LGNN (Meng et al. 2020), TP-LSD (Huang et al. 2020b), LETR (Xu et al. 2021), L-CNN (Zhou, Qi, and Ma 2019), HAWP (Xue et al. 2020) and M-LSD (Gu et al. 2022). All the methods are learning-based methods except the classical LSD. The comparison results are listed in Table 2. RANK-LETR outperforms existing Transformer-based and CNN-based methods in prediction accuracy, particularly on the Wireframe dataset. Additionally, our method can converge with fewer training epochs compared to previous Transformer-based models. We also show the accuracy curves for line segment detection over the time of training in Fig.4. Our method reaches a high level of accuracy after just 60 epochs, which also proves that our method has a fast and stable convergence.

Visual examples of line segment detection results of two Transformer-based methods including LETR and ours on outdoor and indoor scenes in the Wireframe dataset are shown in Fig. 3. It is shown that our method can produce more accurate and complete detection result, especially in the detection of short line segments and in areas with high line density.

### Ablation Study

To verify the effectiveness of components and find the influence of the hyperparameters in our proposed method, we conduct an ablation and parameter study of our network architecture. The results are summarized in Table 3. The ablation study is conducted on the Wireframe dataset and the




Figure 4: The accuracy curves for line segment detection over the time of training. We observed that our method reaches a high level of accuracy after just 60 epochs.

$sAP$  and  $sF$  results are reported.

We employed the combination of ResNet50 and Transformer Encoder and Decoder from the classical LETR method as our baseline approach. We constructed a foundational version of our modeling using a ResNet50 network and a 6-layer Deformable Transformer Encoder, where feature maps at resolutions of  $128 \times 128$ ,  $64 \times 64$ , and  $32 \times 32$  are input from the feature extraction backbone into the Transformer network for processing. Based on this architecture, we can use matched predicting strategy in supervising and find it can significantly improve the detection performance. On this basis, we sequentially added a re-ranking module and a rotation augmentation module, both of which further enhanced performance. Additionally, we compared the results of line segment prediction across different resolution feature maps, concluding that line segment prediction on high-resolution feature maps can provides a higher accuracy.

Matched Predicting	Rotation augmentation	Re-ranking	Resolution	Hidden Dim	Referring Points	Wireframe			
						sAP <sup>5</sup>	sAP <sup>10</sup>	sF <sup>5</sup>	sF <sup>10</sup>
-	-	-	-	-	-	50.3	58.5	56.8	61.5
✓	-	-	128	128	4	60.2	65.9	64.9	66.2
✓	✓	-	128	128	4	61.7	67.4	63.0	66.2
✓	✓	✓	64	128	4	63.4	68.3	62.7	65.5
✓	✓	✓	128	128	4	<b>65.0</b>	<b>69.7</b>	<b>64.2</b>	<b>66.9</b>
✓	✓	✓	128	64	4	62.9	68.3	62.7	65.8
✓	✓	✓	128	256	4	<b>65.1</b>	<b>69.8</b>	<b>64.2</b>	<b>66.8</b>
✓	✓	✓	128	128	2	64.4	69.4	63.7	66.6
✓	✓	✓	128	128	8	64.6	69.6	63.8	66.6

Table 3: Ablation and parameter study of our method on the Wireframe (Huang et al. 2018) dataset. We first construct a baseline method according to our modeling, and then gradually add different components. Experimental results show that applying matched-predicting strategy, re-ranking module, rotation augmentation and predicting on high-resolution feature maps can all boost the performance of our approach. We also select the number of hidden dim and referring points according to the parameter study on them.

Method	sAP <sup>5</sup>	sAP <sup>10</sup>	sF <sup>5</sup>	sF <sup>10</sup>
w/o Ranking Loss	26.5	29.2	36.9	38.9
with Ranking Loss	<b>27.6</b>	<b>30.1</b>	<b>37.9</b>	<b>39.7</b>

Table 4: Line segment detection accuracy on YUD with and without ranking loss. We observe that the ranking loss can improve the generalization capability of the proposed method when trained and evaluated on different datasets.

Moreover, we vary the number of referring points and the hidden dim of the Transformer architecture to show the influence of these hyperparameters. We find the hidden dim may be appropriately set to more than 128 since a smaller one will lead to a lower accuracy. Changing the number of referring points has little impact on detection performance. In the ablation studies above, except for the baseline method that required 240 epochs to ensure convergence, all other methods were trained for 120 epochs, demonstrating the improved convergence capabilities of our proposed strategy. The saliency maps generated from some feature maps are also presented in Fig.5 to exhibit the learning effect. We also demonstrate the effectiveness of the ranking loss in the Table 4. We found that using the ranking loss can improve the performance of the model when trained on the Wireframe dataset and tested on the YUD dataset. This indicates that the ranking loss can enhance the generalization performance of the model.

While our method generally achieves the highest detection accuracy, it has a limitation when the centroids of two line segments are very close. In such cases, our method can only predict one of the segments. Although we find that nearby grid points may predict the corresponding line segment instead of the grid where the centroid lies, this remains a theoretical limitation of our approach. Allowing each feature point to predict more than one, e.g., two line segments could potentially mitigate this issue, which we leave as future work.




Figure 5: Saliency maps generated from score maps (left), edge maps (middle) and endpoint maps (right).

## Conclusion

We identify a critical factor limiting the performance of proposal-based line segment detection algorithms: the confidence scores may fail to accurately reflect the quality of the predicted line segments. To tackle this issue, we propose three novel techniques to enhance Transformer-based line segment detection methods. First, we introduce a simple and efficient line re-ranking module that optimizes the confidence scores of lines using learnable geometric information, such as edges, endpoints, and line lengths. This module is more interpretable and allows for flexible weighting across different scenarios. We also present a matched prediction strategy, wherein each feature point is responsible for detecting the line segment whose centroid is closest to it. Furthermore, we propose a line segment ranking loss to make feature points predict higher confidence scores for higher-quality predicted line segments during training. Building on these techniques, we developed a novel line segment detection model named RANK-LETR, which outperforms existing Transformer-based and CNN-based methods while requiring fewer training epochs than previous Transformer-based models.

## References

- Akinlar, C.; and Topal, C. 2011. EDLines: A real-time line segment detector with a false detection control. *Pattern Recognition Letters*, 32(13): 1633–1642.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*, 213–229. Springer.
- Cetinkaya, B.; Kalkan, S.; and Akbas, E. 2024. RankED: Addressing Imbalance and Uncertainty in Edge Detection Using Ranking-based Losses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 3239–3249.
- Chen, K.; Li, J.; Lin, W.; See, J.; Wang, J.; Duan, L.; Chen, Z.; He, C.; and Zou, J. 2019. Towards accurate one-stage object detection with ap-loss. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5119–5127.
- Chen, X.; Yan, B.; Zhu, J.; Wang, D.; Yang, X.; and Lu, H. 2021. Transformer tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8126–8135.
- Denis, P.; Elder, J. H.; and Estrada, F. J. 2008. Efficient edge-based methods for estimating manhattan frames in urban imagery. In *European conference on computer vision*, 197–210. Springer.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Furukawa, Y.; and Shinagawa, Y. 2003. Accurate and robust line segment extraction by analyzing distribution around peaks in Hough space. *Computer vision and image understanding*, 92(1): 1–25.
- Gomez-Ojeda, R.; Moreno, F.-A.; Zuniga-Noël, D.; Scaramuzza, D.; and Gonzalez-Jimenez, J. 2019. PL-SLAM: A stereo SLAM system through the combination of points and line segments. *IEEE Transactions on Robotics*, 35(3): 734–746.
- Gu, G.; Ko, B.; Go, S.; Lee, S.-H.; Lee, J.; and Shin, M. 2022. Towards light-weight and real-time line segment detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 726–734.
- Guil, N.; Villalba, J.; and Zapata, E. L. 1995. A fast Hough transform for segment detection. *IEEE transactions on image processing*, 4(11): 1541–1548.
- Hofer, M.; Maurer, M.; and Bischof, H. 2017. Efficient 3D scene abstraction using line segments. *Computer Vision and Image Understanding*, 157: 167–178.
- Huang, K.; Wang, Y.; Zhou, Z.; Ding, T.; Gao, S.; and Ma, Y. 2018. Learning to parse wireframes in images of man-made environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 626–635.
- Huang, L.; Tan, J.; Liu, J.; and Yuan, J. 2020a. Hand-Transformer: Non-Autoregressive Structured Modeling for 3D Hand Pose Estimation. In *European Conference on Computer Vision*, 17–33. Springer.
- Huang, S.; Qin, F.; Xiong, P.; Ding, N.; He, Y.; and Liu, X. 2020b. TP-LSD: Tri-points based line segment detector. In *European Conference on Computer Vision*, 770–785. Springer.
- Langlois, P.-A.; Boulch, A.; and Marlet, R. 2019. Surface reconstruction from 3d line segments. In *2019 International Conference on 3D Vision (3DV)*, 553–563. IEEE.
- Li, C.; Liu, P.; and Zhou, Y. 2012. A strategy on corresponding line segment recognition for 3D-reconstruction. In *Future Control and Automation: Proceedings of the 2nd International Conference on Future Control and Automation (ICFCA 2012)-Volume 1*, 77–85. Springer.
- Li, F.; Zhang, H.; Liu, S.; Guo, J.; Ni, L. M.; and Zhang, L. 2022. Dn-detr: Accelerate detr training by introducing query denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13619–13627.
- Lin, Y.; Pinteá, S. L.; and van Gemert, J. C. 2020. Deep hough-transform line priors. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXII 16*, 323–340. Springer.
- Lu, X.; Yao, J.; Li, K.; and Li, L. 2015. Cannylines: A parameter-free line segment detector. In *2015 IEEE International Conference on Image Processing (ICIP)*, 507–511. IEEE.
- Meng, Q.; Zhang, J.; Hu, Q.; He, X.; and Yu, J. 2020. LGNN: A Context-aware Line Segment Detector. In *Proceedings of the 28th ACM International Conference on Multimedia*, 4364–4372.
- Nakano, G. 2021. Camera calibration using parallel line segments. In *2020 25th International Conference on Pattern Recognition (ICPR)*, 1505–1512. IEEE.
- Oksuz, K.; Cam, B. C.; Akbas, E.; and Kalkan, S. 2021. Rank & sort loss for object detection and instance segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3009–3018.
- Tang, X.-s.; Xie, X.; Hao, K.; Li, D.; and Zhao, M. 2022. A line-segment-based non-maximum suppression method for accurate object detection. *Knowledge-Based Systems*, 251: 108885.
- Tong, X.; Peng, S.; Guo, Y.; and Huang, X. 2024. End-to-End Real-Time Vanishing Point Detection with Transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 5243–5251.
- Tong, X.; Ying, X.; Shi, Y.; Wang, R.; and Yang, J. 2022. Transformer Based Line Segment Classifier with Image Context for Real-Time Vanishing Point Detection in Manhattan World. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6093–6102.
- Vakhitov, A.; and Lempitsky, V. 2019. Learnable line segment descriptor for visual SLAM. *IEEE Access*, 7: 39923–39934.
- Von Gioi, R. G.; Jakubowicz, J.; Morel, J.-M.; and Randall, G. 2008. LSD: A fast line segment detector with a false detection control. *IEEE transactions on pattern analysis and machine intelligence*, 32(4): 722–732.

Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34: 12077–12090.

Xu, Y.; Xu, W.; Cheung, D.; and Tu, Z. 2021. Line segment detection using transformers without edges. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4257–4266.

Xu, Z.; Shin, B.-S.; and Klette, R. 2015. Closed form line-segment extraction using the Hough transform. *Pattern Recognition*, 48(12): 4012–4023.

Xue, N.; Bai, S.; Wang, F.; Xia, G.-S.; Wu, T.; and Zhang, L. 2019. Learning attraction field representation for robust line segment detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1595–1603.

Xue, N.; Wu, T.; Bai, S.; Wang, F.; Xia, G.-S.; Zhang, L.; and Torr, P. H. 2020. Holistically-attracted wireframe parsing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2788–2797.

Yavuz, F.; Cam, B. C.; Dogan, A. H.; Oksuz, K.; Akbas, E.; and Kalkan, S. 2024. Bucketed Ranking-based Losses for Efficient Training of Object Detectors. In *European Conference on Computer Vision*. Springer.

Zavala, J. G. N.; and Martinez-Carranza, J. 2022. Depth Estimation from a Single Image Using Line Segments only. In *Ibero-American Conference on Artificial Intelligence*, 331–341. Springer.

Zhang, Y.; Zhou, L.; Liu, H.; and Shang, Y. 2016. A flexible online camera calibration using line segments. *Journal of Sensors*, 2016(1): 2802343.

Zhang, Z.; Li, Z.; Bi, N.; Zheng, J.; Wang, J.; Huang, K.; Luo, W.; Xu, Y.; and Gao, S. 2019. Ppgnet: Learning point-pair graph for line segment detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7105–7114.

Zhou, Y.; Qi, H.; and Ma, Y. 2019. End-to-end wireframe parsing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 962–971.

Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.