

# SOFI: Multi-Scale Deformable Transformer for Camera Calibration with Enhanced Line Queries

Sebastian Janampa Student  
sebasjr1966@unm.edu

Marios Pattichis  
pattichi@unm.edu

Department of Electrical & Computer  
Engineering  
The University of New Mexico  
New Mexico, USA

---

## Abstract

Camera calibration consists of estimating camera parameters such as the zenith vanishing point and horizon line. Estimating the camera parameters allows other tasks like 3D rendering, artificial reality effects, and object insertion in an image. Transformer-based models have provided promising results; however, they lack cross-scale interaction. In this work, we introduce *multi-Scale defOrmable transFormer for camera calibration with enhanced line queries*, SOFI. SOFI improves the line queries used in CTRL-C and MSCC by using both line content and line geometric features. Moreover, SOFI's line queries allow transformer models to adopt the multi-scale deformable attention mechanism to promote cross-scale interaction between the feature maps produced by the backbone. SOFI outperforms existing methods on the *Google Street View*, *Horizon Line in the Wild*, and *Holocity* datasets while keeping a competitive inference speed. Code is available at: <https://github.com/SebastianJanampa/SOFI>

## 1 Introduction

Camera calibration is a fundamental task in computer vision, which enables higher-level applications such as 3D scene renderization, image rectification, and metrology. Estimating camera (intrinsic and extrinsic) parameters from a single-view image involves the development of methods that process the image's perspective distortions. Here, note that an example of perspective distortion is the intersection of lines on an image plane despite being parallel in the real world.

Traditional methods involve using known calibration patterns (e.g., chessboard images). These methods include taking multiple pictures of the 3D environment with the calibration pattern at different locations and angles. Then, corners are extracted from each image and matched using a matching technique such as Random Sample Consensus (RANSAC) [1] to compute the camera parameters. Traditional methods perform well, but require processing images of calibrated patterns and do not allow automated calibration from a single image.

Alternatively, camera calibration can be performed based on the relationship between camera parameters and geometric cues, such as vanishing points and horizon lines. We summarize these approaches based on three main steps: line segment detection, line clustering,

and vanishing point estimation from each cluster. Unfortunately, such methods depend on man-made scenes (e.g., buildings). Their performance drops significantly on natural images or images taken in the wild.

Deep learning methods rely on both geometric cues and image features. In earlier research, Deepfocal [23] proposed a deep learning model to estimate the focal length of a camera based on any image taken "in the wild". Workman *et al.* [24] proposed a convolutional neural network (CNN) to predict the horizon line. Deepcalib [2] proposed a CNN model to predict the intrinsic (horizon line, focal length, and distortion) camera parameters. Hold-Geoffroy *et al.* proposed a model that estimates both intrinsic and extrinsic at the same time. Authors in [10, 28] use geometric priors (lines) to improve estimation. Overall, deep learning methods based on convolutional layers can only capture dependencies over the support of the convolution kernels. Here, we note that camera calibration requires extracting information from different regions of the image that are not necessarily near each other.

More recently, Camera Calibration Transformer with Line-Classification (CTRL-C) [11] and Multi-Scale Camera Calibration (MSCC) [17] introduced transformer-based [8, 20] models that use the attention mechanism to extract long-term dependencies over the input images. MSCC used a coarse-to-fine technique that extended CTRL-C output queries for coarse estimation. Then, it used a second transformer model fed by a larger feature map with twice the spatial resolution of the first map. However, the use of a larger feature map increases inference time. At the same time, in MSCC, line geometric information is not seen by higher layers of the decoder transformer.

We propose a new neural network approach called *multi-Scale defOrmable transFormer for camera calibratIon with enhanced line queries* (SOFI). Our model uses deformable attention to extract information from higher-resolution feature maps. Plus, SOFI uses line segment geometric information as input for each encoder layer.

Our contributions are summarized as follows:

- We propose a new initialization for lines' queries which not only improves transformer-based models for camera calibration, but it also allows us to use the deformable attention mechanism.
- We reformulate the line classification functions. Additionally, we update the loss functions' coefficients, giving more importance to the camera parameters losses than to the line classification losses.
- We propose a novel framework that achieves new state-of-the-art results in in- and out-of-distribution datasets.

## 2 Background

### 2.1 Camera Calibration

Camera calibration methods are used for estimating camera parameters. Earlier methods for camera calibration are based on the use of vanishing points (VPs) and the horizon line. Schafalitzky and Zimmerman [25] proposed an automatic VP detection model by grouping geometric features. Tretyak [19] presented a parsing framework to estimate VPS by geometrically analyzing low-level geometric features to estimate the zenith VP and the horizon line. However, to be effective, these approaches require images of human-made structures

(e.g., buildings), rich in geometric content that satisfy the Manhattan or Atlanta environment assumptions.

On the other hand, early deep learning models directly estimated the camera parameters by extracting geometric cues using convolutional layers [23, 24]. To improve results, deep learning models used line segments detected by LSD [21] algorithm to better understand the scene. However, these approaches require post-processing steps and have complicated architectures for integrating the line information into the deep learning model.

Lee *et al.* [10] introduced Camera calibration TRansformer with Line-Classification (CTRL-C), the first end-to-end transformer model for camera calibration. The network used a CNN backbone to extract image features and feed them into an encoder-decoder transformer model. It also utilized line information to achieve better results. Nevertheless, in CTRL-C, the input was restricted to the lowest-resolution feature map, while higher resolution feature maps were ignored. Song *et al.* [17] observed that low-resolution (high-level) feature maps contain global information about an image, while higher-resolution (low-level) features capture image details. The authors [17] propose MSCC, a multi-scale transformer model for camera calibration. MSCC uses higher-level feature maps to retrieve information about image structures and low-level feature maps to get finer details about vanishing points and horizon lines. Despite achieving new state-of-the-art (SOTA) results, standard transformers models are not time-efficient for processing low-level features because of high complexity associated with the encoder stage.

## 2.2 DETR-based methods

Since Carion *et al.* introduced DEtection TRansformer (DETR), many methods have been proposed to improve the DETR (e.g., see [18]). In this section, we provide a short description of the DETR model and how it is used in CTRL-C [10] and MSCC [17]. Then, we describe the deformable attention mechanism proposed in deformable DETR [30]. We end the section by discussing about the query formulation used in end-to-end transformer models.

### 2.2.1 Encoder in Detection Transformer for Camera Calibration

DEtection TRansformer [9] is the first end-to-end transformer model that discards post-processing steps. It consists of a backbone, an encoder and a decoder. The backbone produces a set of feature maps where only the highest-level feature map goes to the encoder for feature enhancing. This feature encoding consists on intra-scale processing using an attention mechanism [20] defined as:

$$\text{Attn} = \text{Softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (1)$$

where  $Q$ ,  $K$ , and  $V$  represent the queries, keys, and values, respectively. For the encoder part, we have that  $Q = K = V = f \in \mathbb{R}^{(HW) \times d}$  where  $f$  is the given feature map of height  $H$ , width  $W$ , and number of channels  $d$ .

CTRL-C [10] and MSCC [17] follows the previously described feature map enhancing method. The difference between the two models is that CTRL-C applies it to only the highest-level feature map, and MSCC to the two highest-level feature maps using a different encoder for each feature map. Then, each enhanced feature map is passed to a decoder to estimate the camera parameters.

In the encoder, CTRL-C estimates the camera parameters using a global feature map. Although the pixels of the feature map are rich in information, the spatial size is too small. On the other hand, MSCC uses CTRL-C output queries as pre-initialized queries to process them using a second decoder block with the second highest-level feature map. However, the global information is lost through each layer of the second decoder block. We solve this issue by using the deformable attention mechanism to promote cross-scale processing while keep a good complexity time for the encoder block.

### 2.2.2 Deformable Attention Mechanism

The deformable attention mechanism is proposed in DeformableDETR [60] to allow DETR to use high-resolution feature maps without compromising inference time. The idea consists on compute  $K$  sampling points as well as an attention matrix from the feature map for each query  $q \in Q$  such as  $K \ll HW$ , reducing the complexity time from  $O(H^2W^2d)$  to  $O(KHWd)$ <sup>1</sup>. For a single-head deformable attention mechanism, the equation is

$$\text{DeformAttn}(z_q, p_q, x) = \sum_{k=1}^K A_{qk} \cdot x(p_q + \Delta p_{qk}) \quad (2)$$

where  $k$  indexes the sampling keys. The  $k^{\text{th}}$  sampling key is computed as  $p_q + \Delta p_{qk}$  where  $\Delta p_{qk}$  is the sampling offset and  $p_q$  is the 2d-reference point for query  $q$ .  $A_{qk}$  is the  $k^{\text{th}}$  row of the attention weight  $A_q \in \mathbb{R}^{K \times d}$ . The weights of  $A_{qk}$  satisfy  $\sum_{k=1}^K A_{qk} = 1$ .

In addition to the time complexity reduction, DeformAttn has a variation MSDeformAttn, where MS stands for Multi-Scale, that promotes cross-scale interaction. Similar to DeformAttn, given a set of  $L$  feature maps  $\{x^l\}_{l=1}^L$ , MSDeformAttn samples  $K$  sampling point for each  $x^l$ . The mathematical representation for a single-head multi-scale deformable attention is

$$\sum_{l=1}^L \sum_{k=1}^K A_{lqk} \cdot x^l(\phi_l(\hat{p}_q) + \Delta p_{lqk}) \quad (3)$$

where  $x^l \in \mathbb{R}^{H_l \times W_l \times d}$ .  $\hat{p}_q$  is the normalized 2d coordinates whose values lie in the range of  $[0, 1]$ , and that is re-scaled to dimensions of the feature map of the  $l^{\text{th}}$  level by the function  $\phi_l$ . Like eq. (2), the attention weight  $A_{lqk}$  satisfies  $\sum_{l=1}^L \sum_{k=1}^K A_{lqk} = 1$ .

For the encoder, a query  $q$  corresponds to a pixel in  $x^l$  meaning that the pixel will interact with  $K$  points from  $x_l$  (intra-scale interaction), and  $K$  points from the each of the remaining  $(L - 1)$  feature maps (cross-scale interaction). This allows the propagation of global information to low-level feature maps.

### 2.2.3 Queries in Decoder Transformers

In this section, we summarize prior related research on modifying the query inputs for DETR-based models. DETR [8] uses learnable vectors to provide positional constraints. However, this works assuming that the learnable vectors contain coordinate information. Dynamic Anchor Boxes (DAB)-DETR [13] and Anchor-DETR [22] proposed anchor boxes and points to extract positional information and improve decoder performance.

In CTRL-C [14] and MSCC [17], line content queries were initialized with the line geometric information, while the positional queries were initialized with zero vectors. The

<sup>1</sup>Refer to [60] for the full derivations of the time complexities.

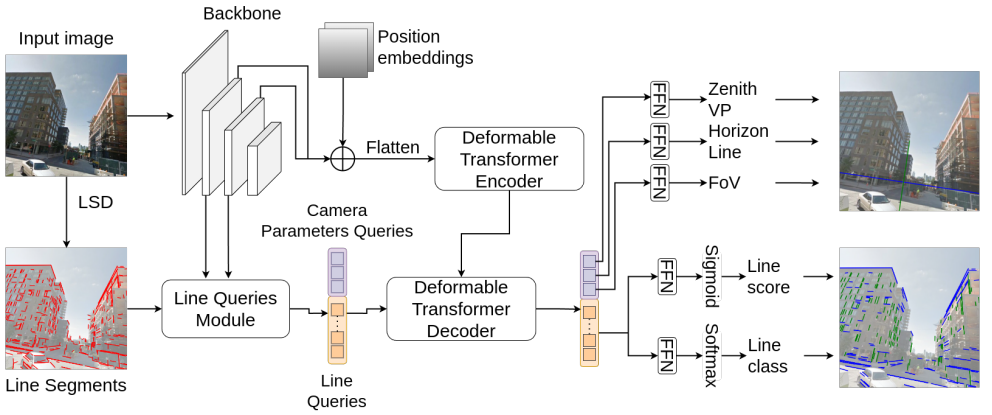


Figure 1: Model overview. Our proposed network uses ResNet50 [6] as backbone to extract feature maps from stage 2 and 3 which are then fed to the deformable transformer encoder.

approaches share the first decoder layer as done in the vanilla DETR [6] model. Unfortunately, in CTRL-C and MSCC, line geometric information is lost through each layer because the positional query is zero. Also, their use of positional queries with zero vectors makes it impossible to use a deformable attention mechanism that requires the use of positional queries to extract reference points.

## 3 Methodology

### 3.1 Model Overview

We present the SOFI architecture in Fig. 1. The input image is processed using ResNet50 [6] to generate multi-scale feature maps, which are then used as encoder input and to prepare the line content queries. The encoder enhances the backbone feature maps and passes them to the decoder. The decoder outputs are processed using Feed-Forward Networks (FFNs), which produce estimates of the camera parameters (see top-right of the figure) and estimate the line segment in the image (see bottom-right of the figure).

In what follows, we provide a detailed explanation of our modified line queries and the line classification method. Our deformable encoder-decoder transformer network is adopted from Deformable DETR [60]. For our SOFI network, we use 32 sampling offsets in the deformable encoder, and 8 sampling offsets in the deformable decoder (see section 4.1 from [60] for information about the deformable attention mechanism and the sampling offsets)

#### 3.1.1 Line Queries Module

Our Queries consist of content-based components  $q_{con}$  and position-based components  $q_{pos}$ .

We follow the same procedure for position queries as in [14, 17]. Given the two endpoints of a line, we compute the homogeneous line representation  $l = [a, b, c]$  where  $c$  is the line offset and  $a/b$  is the line slope. We also remove directional ambiguities ( $l$  and  $-l$  represents the same line) using  $s = [a^2, ab, b^2, bc, ac, c^2]$ .

Next, we define our line position-based component using:  $q_{pos} = \text{FCL}(s)$  where FCL is a fully-connected layer that maps  $s$  to a  $d$ -dimensional vector, where  $d$  represents the input size of the decoder.

Similarly, we define our content queries using  $q_{con} = \text{FCL}(\text{LS}(\mathbf{p}, \mathbf{F}))$  where  $\mathbf{p}$  represents the lines' endpoints,  $\mathbf{F}$  represents the chosen feature maps, and LS (for line sampling) generates  $N = 16$  uniformly sampled points between the endpoints.

### 3.1.2 Line Classification

We classify each detected line segment into either (i) horizontal line, (ii) vertical line, or (iii) other. In addition, for horizontal and vertical line classifications, we compute a confidence score that reflects the importance of using the detected line for camera calibration. We describe how to setup the classification loss and confidence loss functions in the section 3.2. Here, we describe our line classifier and the regression model used for estimating the confidence level.

We use a fully-connected layer followed by a sigmoid activation function for implementing line classification. The classifier input comes from the output of the deformable decoder. More precisely, our classifier is given by  $\mathcal{F}_{class} : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}^{N \times 3}$  where  $N = 512$  represents the number of detected lines and  $d = 256$  represents the dimension of the decoder vector output for each line.

Similarly, we use a fully-connected layer followed by sigmoid activation function for estimating the confidence score. Thus, our regression network is  $\mathcal{F}_{score} : \mathbb{R}^{N \times d} \rightarrow \mathbb{R}^{N \times 1}$ . Here, we have the same input as for line classification. The output represents the confidence score for each one of our  $N$  lines.

## 3.2 Loss function

As shown in Fig. 1, we need to design loss functions for each output. We use: cosine distance for Zenith VP, and  $L_1$  loss for Horizon Line and field of view (FoV) (see the Supplementary Material). We use the focal loss for the estimated classification probability of each class and the estimated confidence level as given by [10]:

$$FL(q) = -\frac{1}{N} \sum_{i=0}^N (\alpha_1 (1 - q^{(i)})^\gamma \log q^{(i)} + \alpha_2 (q^{(i)})^\gamma \log (1 - q^{(i)})), \quad (4)$$

where  $q$  is the confidence score or the classification probability,  $\alpha_1$  and  $\alpha_2$  are the weighted coefficients that are determined by distances to the vanishing points as given in section 3.5 of [10], and  $\gamma = 2$ .

Our overall loss  $\mathcal{L}$  is defined using:

$$\mathcal{L} = 5\mathcal{L}_{zvp} + 5\mathcal{L}_{hl} + 5\mathcal{L}_{FoV} + \mathcal{L}_{score} + \mathcal{L}_{class} \quad (5)$$

where the  $\mathcal{L}_{zvp}$ ,  $\mathcal{L}_{hl}$  and  $\mathcal{L}_{FoV}$  define the loss functions of the camera parameters, while  $\mathcal{L}_{class}$  and  $\mathcal{L}_{score}$  are the loss functions for the confidence and the line class probability that use Eq. (4). Here, we place more importance to estimating camera parameters than line classification.

Model	Up (°) ↓		Pitch (°) ↓		Roll (°) ↓		FoV (°) ↓	
	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.
Google Street View [10]								
Upright [8]	3.05	1.92	2.90	1.80	6.19	0.43	9.47	4.42
DeepHorizon [12]	3.58	3.01	2.76	2.12	1.78	1.67	-	-
Perceptual [8]	2.73	2.13	2.39	1.78	0.96	0.66	4.61	3.89
UprightNet [12]	28.20	26.10	26.56	24.56	6.22	4.33	-	-
GPNet [11]	2.12	1.61	1.92	1.38	0.75	0.47	3.59	2.72
CTRL-C [10]	1.71	1.43	1.52	<b>1.20</b>	0.57	0.46	3.38	2.64
MSCC [10]	1.75	<b>1.42</b>	1.56	1.24	0.58	0.46	<b>3.04</b>	<b>2.29</b>
SOFI (ours)	<b>1.64</b>	1.44	<b>1.51</b>	1.28	<b>0.54</b>	<b>0.43</b>	3.09	2.79
Holicity [14]								
DeepHorizon* [12]	7.82	3.99	6.10	2.73	3.97	2.67	-	-
Perceptual* [8]	7.37	3.29	6.32	2.86	3.10	1.82	-	-
GPNet* [11]	4.17	<b>1.73</b>	<b>1.46</b>	<b>0.74</b>	1.36	0.95	-	-
CTRL-C [10]	2.66	2.19	2.26	1.78	1.09	0.77	12.41	11.59
MSCC [10]	2.28	1.88	1.87	1.43	<b>1.08</b>	<b>0.81</b>	13.60	12.20
SOFI (ours)	<b>2.23</b>	1.82	1.75	1.31	1.16	0.85	<b>11.47</b>	<b>11.25</b>

Table 1: Results of camera calibration parameters on testing datasets. The \* is used to mark models that were trained using the SUN360 [16] dataset.

## 4 Results

### 4.1 Dataset

**Google Street View:** In our experiments, we train all models with 12679 images from the Google Street View dataset [10]. The dataset consists of images of streets, buildings and landmarks that satisfy the Manhattan World assumption (three orthogonal vanishing points). Additionally, this dataset contains 535 images for validation, and 1333 for testing.

**Horizon Line in the Wild:** Horizon Line in the Wild (HLW) dataset [14] only provides ground-truth for the horizon lines. We use 2019 images from this dataset to test horizon line detection in unseen data. Here, by unseen data, we refer to the fact that none of the models are trained on this dataset.

**Holicity:** We test SOFI on the holicity dataset [14] to estimate all camera parameters. The dataset consists of 2024 images. No model is trained on this dataset.

### 4.2 Implementation Details

We train all the models using a single Nvidia RTX A5500 GPU for 30 epochs. For SOFI, we use a learning rate of  $2 \times 10^{-4}$  for the first 20 epochs. For the last 10 epochs, we set the learning rate to  $2 \times 10^{-5}$ . We use AdamW for optimization with weight decay of  $10^{-4}$ .

We train MSCC using the recipe from LETR [17] due to the lack of training information in the original paper. First, we train the coarse stage for 30 epochs. Second, we train the coarse and fine stages for another 30 epochs. For MSCC and CTRL-C, we use a learning of  $10^{-4}$ , and we decrease the value to  $10^{-5}$  at epoch 21.

Model	Google Street View [1]			Horizon Line in the Wild [2]			Holicity [3]		
	@ 0.10	@ 0.15	@ 0.25	@ 0.10	@ 0.15	@ 0.25	@ 0.10	@ 0.15	@ 0.25
DeepHorizon* [2]	-	-	74.25	-	-	45.63	-	-	70.13
Perceptual* [1]	-	-	80.40	-	-	38.29	-	-	70.80
GPNet* [1]	-	-	83.12	-	-	48.90	-	-	81.72
CTRL-C	69.49	78.92	87.16	24.04	33.56	46.37	38.84	55.13	72.31
MSCC	<b>70.39</b>	79.59	87.63	24.85	34.44	47.28	49.71	63.60	77.43
SOFI (ours)	70.32	<b>79.84</b>	<b>87.87</b>	<b>27.93</b>	<b>37.55</b>	<b>49.69</b>	<b>59.83</b>	<b>72.05</b>	<b>82.96</b>

Table 2: AUC percentages (out of 100%) for horizon line errors on testing datasets. The \* is used to mark models that were trained using the SUN360 [26] dataset. @ 0.10, @ 0.20, and @ 0.25 refer to the area under the curve from zero to Error=0.10, 0.20, and 0.25 as shown in Fig. 3.

Model	Google Street View [1]	Horizon Line in the Wild [2]	Holicity [3]
CTRL-C [1]	<b>22.6</b>	<b>19.0</b>	<b>25.9</b>
MSCC [1]	18.0	13.2	18.4
SOFI (ours)	21.6	17.4	23.6

Table 3: Inference speed comparisons for transformers-based models with batch size of 1. Results are shown in terms of frames per second. Refer to section 4.2 for hardware description.

### 4.3 Comparisons

We compare SOFI against Upright [1], DeepHorizon [2], UprightNet [5], GPNet [1], CTRL-C [1], and MSCC [1]. DeepHorizon, Perceptual, and GPTNet were trained on the SUN360 [26] dataset, which is no longer available (due to license issues) for the Holicity dataset. For evaluation metrics, we use the up-vector, two angles of the rotation matrix (pitch and roll), the field of view (FoV), and the area under the curve (AUC) percentage error for the horizon line. For all the experiments, we consider a camera model with no skew (square pixels), and yaw = 0 [1, 1, 1] (the third angle in the camera rotation matrix). For information about the estimation of the up vector, the roll, and pitch using the zenith VP and the FoV, see section 3.1 in [1].

We provide comparative results for the camera parameters and the horizon line in Table 1 and Table 2, respectively. SOFI provides the best results on the Google Street View dataset but for a small margin of difference. We believe this happens because the dataset is too relatively easy to learn or the training and testing sets are very similar. On the other hand, SOFI increases the margin of difference for out-of-distribution datasets (Holicity and HWL datasets). We find the most significant improvement is in the horizon line detection, where there is an improvement of 5 points with respect to MSCC, as shown in Tab. 2 and Fig. 3. We want to clarify that the FoV error for the Holicity dataset is large because the sampling range of this dataset is bigger than the one in the training dataset.

We provide horizon line estimation examples for transformer-based models in Fig. 2. We note that SOFI provides consistently accurate estimates in all of the examples. In contrast, we can see that CTRL-C and MSCC can be very inaccurate for the unseen datasets of the bottom two rows.

We present inference speed results in table 3. Despite the fact that SOFI uses multiple feature maps, at higher resolutions than MSCC and CTRL-C, our inference speed remains very competitive. SOFI runs faster than MSCC. SOFI performs nearly as fast as CTRL-C which only uses a single low-resolution feature map.





Ground Truth CTRL-C MSCC SOFI (Ours)

Figure 2: Examples of horizon line estimation on the Google Street View [10] test set (top row), the Horizon Line in the Wild [24] test set (middle row), and the Holicity [29] test set (bottom row).




Figure 3: Cumulative distribution error for the horizon line on Horizon Line in the Wild [24] and Holicity test set [29].

#### 4.4 Ablation Study

We estimate the camera parameters using the original Deformable DETR (with four feature maps) [60] as a baseline. We modify it to use only the line geometry information to evaluate against CTRL-C [10]. As shown in Tab. 4, Deformable-DETR beats CTRL-C in the Google Street View dataset, but it does not on the Holicity dataset. This performance occurs because of the line query formulation CTRL-C has, which does not allow the estimation of proper sampling points. We refer the reader to the supplementary material for more information about this problem. There, we conducted experiments using only CTRL-C to validate the importance of choosing good coefficients for each loss and our new line queries that use geometric and content information.

Model	Up (°) ↓	Pitch (°) ↓	Roll (°) ↓	FoV (°) ↓	AUC (%) ↑	FPS ↑
	Google Street View [10]					
CTRL-C [10]	1.71	1.52	0.57	3.38	87.16	22.6
Deformable-DETR	1.68	1.51	0.54	3.13	87.57	19.6
SOFI* <sup>†</sup>	1.69	1.49	0.61	3.03	87.48	22.1
SOFI <sup>†</sup>	1.66	1.49	0.55	3.11	87.10	21.4
SOFI	1.64	1.51	0.54	3.09	87.87	21.6
SOFI-B	1.58	1.47	0.59	2.94	89.12	18.2
	Holicity [10]					
CTRL-C [10]	2.66	2.26	1.09	12.41	72.31	25.9
Deformable-DETR	2.78	2.45	1.03	10.93	71.56	21.0
SOFI* <sup>†</sup>	2.52	2.14	1.07	12.02	83.18	25.2
SOFI <sup>†</sup>	2.27	1.81	1.22	12.53	79.82	23.1
SOFI	2.23	1.75	1.16	11.47	82.96	23.6
SOFI-B	2.15	1.70	1.12	11.22	83.21	18.9

Table 4: Ablation study. \* means the model uses the last two feature maps from the backbone. †: means the mode uses the same line classification method as in [10]. SOFI-B uses the last three feature maps. All SOFI variations were trained with 32 and 8 sampling offsets for the encoder and decoder, respectively.

We create SOFI\*<sup>†</sup> to improve the speed and precision. SOFI\*<sup>†</sup> uses the last two feature maps and incorporates the line content information. These modifications not only improved the speed but also improved the metrics, beating CTRL-C on both datasets. Then, we replace the highest-level feature with the third highest-level one. The results are shown in SOFI<sup>†</sup>. Next, we replace the line classification proposed in [10] with ours described in sec. 3.1.2. The results in SOFI prove that our line classification module helps the model to extract richer information about line content.

We also create SOFI-Big (SOFI-B), a modified SOFI network that uses the last three feature maps. As expected, using more feature maps results in better camera parameter estimations, as reflected in the reduced errors for FoV and AUC for the horizon line. Although the significant increment in the FPS for minor accuracy improvements makes SOFI-B unsuitable for applications, it opens new borders to investigate faster cross-scale interaction methods.

## 5 Conclusions

We present a new model that uses deformable attention to produce new SOTA results while operating at competitive inference speed. Our model also introduces a new line query approach that provides better camera calibration estimates than what can be achieved with DETR architectures. Our experiments document significant performance improvements in unseen datasets (Holicity and Horizon line in the Wild datasets).

In future work, we will consider a deeper study of the encoder because of its important role in feature maps enhancing. An important observation from this work is that increasing the number of sampling points in the deformable encoder produces better feature maps. We will consider increasing the number of points without compromising speed. Another direction is to explore different cross-scale interaction mechanism that lead to an acceptable time complexity.

## Acknowledgment

This work was supported in part by the National Science Foundation under Grant 1949230, Grant 1842220, and Grant 1613637

## References

- [1] Google street view images api. URL <https://developers.google.com/maps/>.
- [2] Oleksandr Bogdan, Viktor Eckstein, Francois Rameau, and Jean-Charles Bazin. Deepcalib: a deep learning approach for automatic intrinsic calibration of wide field-of-view cameras. In *Proceedings of the 15th ACM SIGGRAPH European Conference on Visual Media Production*, 2018.
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021.
- [5] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981.
- [6] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [7] Markus Hofinger, Samuel Rota Bulò, Lorenzo Porzi, Arno Knapitsch, Thomas Pock, and Peter Kotschieder. Improving optical flow on a pyramid level. In *European Conference on Computer Vision*, pages 770–786. Springer, 2020.
- [8] Yannick Hold-Geoffroy, Kalyan Sunkavalli, Jonathan Eisenmann, Matthew Fisher, Emiliano Gambaretto, Sunil Hadap, and Jean-François Lalonde. A perceptual measure for deep single image camera calibration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2354–2363, 2018.
- [9] Hyunjoon Lee, Eli Shechtman, Jue Wang, and Seungyong Lee. Automatic upright adjustment of photographs with robust camera calibration. *IEEE transactions on pattern analysis and machine intelligence*, 36(5):833–844, 2013.
- [10] Jinwoo Lee, Minhyuk Sung, Hyunjoon Lee, and Junho Kim. Neural geometric parser for single image camera calibration. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 541–557. Springer, 2020.

- [11] Jinwoo Lee, Hyunsung Go, Hyunjoon Lee, Sunghyun Cho, Minhyuk Sung, and Junho Kim. Ctrl-c: Camera calibration transformer with line-classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16228–16237, 2021.
- [12] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [13] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. DAB-DETR: Dynamic anchor boxes are better queries for DETR. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=oMI9PjOb9Jl>.
- [14] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [15] Frederik Schaffalitzky and Andrew Zisserman. Planar grouping for automatic detection of vanishing lines and points. *Image and Vision Computing*, 18(9):647–658, 2000.
- [16] Tahira Shehzadi, Khurram Azeem Hashmi, Didier Stricker, and Muhammad Zeshan Afzal. Object detection with transformers: A review, 2023.
- [17] Xu Song, Hao Kang, Atsunori Moteki, Genta Suzuki, Yoshie Kobayashi, and Zhiming Tan. Msc: Multi-scale transformers for camera calibration. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 3262–3271, January 2024.
- [18] Xin Tong, Xianghua Ying, Yongjie Shi, Ruibin Wang, and Jinfa Yang. 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 (CVPR)*, pages 6093–6102, June 2022.
- [19] Elena Tretyak, Olga Barinova, Pushmeet Kohli, and Victor Lempitsky. Geometric image parsing in man-made environments. *International Journal of Computer Vision*, 97: 305–321, 2012.
- [20] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL [https://proceedings.neurips.cc/paper\\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf).
- [21] R.G. von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall. LSD: A fast line segment detector with a false detection control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(4):722–732, apr 2010. doi: 10.1109/tpami.2008.300.

- [22] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based detector. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2567–2575, 2022.
- [23] Scott Workman, Connor Greenwell, Menghua Zhai, Ryan Baltenberger, and Nathan Jacobs. Deepfocal: A method for direct focal length estimation. In *2015 IEEE International Conference on Image Processing (ICIP)*, pages 1369–1373, 2015. doi: 10.1109/ICIP.2015.7351024.
- [24] Scott Workman, Menghua Zhai, and Nathan Jacobs. Horizon lines in the wild. In *British Machine Vision Conference (BMVC)*, 2016.
- [25] Wenqi Xian, Zhengqi Li, Matthew Fisher, Jonathan Eisenmann, Eli Shechtman, and Noah Snavely. Uprightnet: geometry-aware camera orientation estimation from single images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9974–9983, 2019.
- [26] Jianxiong Xiao, Krista A Ehinger, Aude Oliva, and Antonio Torralba. Recognizing scene viewpoint using panoramic place representation. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2695–2702. IEEE, 2012.
- [27] 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.
- [28] Menghua Zhai, Scott Workman, and Nathan Jacobs. Detecting vanishing points using global image context in a non-manhattan world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5657–5665, 2016.
- [29] Yichao Zhou, Jingwei Huang, Xili Dai, Linjie Luo, Zhili Chen, and Yi Ma. HoliCity: A city-scale data platform for learning holistic 3D structures. 2020. arXiv:2008.03286 [cs.CV].
- [30] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020.