PBFORMER: CAPTURING COMPLEX SCENE TEXT SHAPE WITH POLYNOMIAL BAND TRANSFORMER

Anonymous authors
Paper under double-blind review

ABSTRACT

We present PBFormer, an efficient yet powerful scene text detector that unifies the transformer with a novel text shape representation **P**olynomial **B**and (PB). The representation has four polynomial curves to fit a text's top, bottom, left, and right sides, which can capture a text with a complex shape by varying polynomial coefficients. PB has appealing features compared with conventional representations: 1) It can model different curvatures with a fixed number of parameters, while polygon-points-based methods need to utilize a different number of points. 2) It can distinguish adjacent or overlapping texts as they have apparent different curve coefficients, while segmentation-based methods suffer from adhesive spatial positions. PBFormer combines the PB with the transformer, which can directly generate smooth text contours sampled from predicted curves without interpolation. To leverage the advantage of PB, PBFormer has a parameter-free cross-scale pixel attention module. The module can enlarge text features and suppress irrelevant areas to benefit from detecting texts with diverse scale variations. Furthermore, PBFormer is trained with a shape-contained loss, which not only enforces the piecewise alignment between the ground truth and the predicted curves but also makes curves' position and shapes consistent with each other. Without bells and whistles about text pre-training, our method is superior to the previous state-of-the-art text detectors on the arbitrary-shaped CTW1500 and Total-Text datasets. Codes will be public.

1 INTRODUCTION

Scene text detection is an active research topic in computer vision and enables many downstream applications such as image/video understanding, visual search, and autonomous driving (Radford et al., 2021; Long et al., 2021; Reddy et al., 2020). However, the task is also challenging. One non-negligible reason is that the text instance can have a complex shape due to the non-uniformity of the text font, skewing from the photograph, and specific art design. Capturing complex text shapes needs to develop effective text representation. State-of-the-art methods roughly tackle this problem with two types of representations. One is the point-based representation, which predicts the points on the image space to control the shape of the points, including the Bezier control points (Liu et al., 2020) and polygon points (Zhang et al., 2021). The other produces segmentation maps. The map can describe the text of various shapes and can benefit from the prediction results at the pixel level (Liao et al., 2020; Zhu et al., 2021b).

Despite the good performance, both types of representation have limitations: 1) Points-based methods suffer from a fixed number of control points (Tang et al., 2022; Zhang et al., 2022). Too few points cannot handle the highly-curved texts, while simply adding points will increase redundancy for most perspective texts. 2) Segmentation-based methods frequently fail in dividing adjacent texts due to ambiguous spatial positions. The produced segmentation map still needs post-processing and often requires extensive training data (Zhu et al., 2021b).

To address these limitations, we propose a novel representation, named **P**olynomial **B**and (PB). In particular, PB consists of four polynomial curves, each of which fits along a text's top, bottom, left, and right sides. It has clear advantages compared with previous text representations. First, the polynomial coefficients of the PB curves are separated in the parameters space even though the two texts are very close in the image space. Example results are shown in Fig. 1a and Fig. 1b. Second, PB

(a) False Contours Crossing by TESTR(Bezier) (b) Clear Contour Separation by PBFormer

Figure 1: **Advantages of PB.** Comparing (a) and (b), PB divides adjacent texts more clearly than Bezier control points. (c) shows the number of output point increase gradually to represent shapes from straight to highly curved. (d) shows varying curve coefficients can handle dynamic shapes with two boundary variables.

is more compact than polygon points. Polynomial curves with very few parameters, such as quadric curves with $K = 3$ shown in Fig. 1(c)(d), can handle various shapes from a straight line to a round curve. Third, there is a direct correspondence between the PB and the text contour. We can directly compare the ground truth contour points with the sampled points from polynomial curves in the training phrase. While either the Bezier-curve-based methods or the segmentation-based methods need to generate the intermediate ground truth from the annotated polygon points, which leads to misalignment with the original annotations.

Witnessing the great success in natural language processing (NLP) (Vaswani et al., 2017), there has been a recent surge of interest in introducing transformers to vision tasks, including scene text detection. The current transformer-based text detectors are with two stages. For example, FewBetter (Tang et al., 2022) adopts a CNN network to segmentation masks, then incorporates a transformer to generate the control points of the polygon or Bezier-Curve. TESTR (Zhang et al., 2022) also decodes the control point in each region of interest after extracting the bounding box proposals. They sacrifice efficiency due to the separated stages and destroy the simplicity of the detection transformer scheme.

We equip the transformer with PB. Instead of adopting the two-stage pipeline, we insert a parameter-free cross-scale pixel attention module between the CNN feature and the transformer encoder-decoder layers. The module enlarges CNN feature maps and generates multi-scale attention maps. After the attention module, we can directly send the per-pixel features to the transformer without inserting a region proposal module or grouping the segmentation results. The transformer decodes each polynomial curve's K coefficients and 2 boundary variables that determine the curve's definition domain. We uniformly sample the points on the predicted curves within the definition domain and compare them with the corresponding points on the ground truth polygon. Such a design supervises the curve piece by piece and can learn the curve shape and range consistently. In summary, the contribution of PBFormer is:

- A novel text representation called the Polynomial Band (PB) is proposed. PB can utilize a fixed number of parameters to capture the text instance with various curvatures. It also excels at distinguishing the spatially close text instances.
- A cross-scale pixel attention module is proposed. The module performs pixel-wise attention across the feature maps with different sizes. It implicitly highlights the text regions and enables the transformer to direct take all pixel-wise features as input.
- We design a shape-constrained loss function. The loss enforces the piece-wise supervision over the predicted curve and consistently optimizes the curve coefficients and definition domains.

Experiments on multi-oriented and curved text detection datasets CTW1500 (Liu et al., 2019) and Total-Text (Chng & Chan, 2017) demonstrate the effectiveness of our approach. Without any pre-training on large-scale text datasets, our method can achieve better results in terms of F-measure. Due to the lightweight network architecture, our method runs real-time and 4.4 $\times$ faster than other open-sourced transformer text detectors.

2 RELATED WORK

Text representation. Conventional text representation can be roughly divided into point-based and segmentation-based methods. With the need to capture more complex text shapes, the point-based methods gradually involve more points for text representation, including 2 points bounding box (Li et al., 2017), 4 points quadrilateral (Li et al., 2019; Sun et al., 2018), 16 points polygon (Zhang et al., 2020) or 30 points polygon (Zhang et al., 2020) changed by texts' length. ABCNet (Liu et al., 2020) calculates Bezier control points (8 points), which is sufficient for most quadrilateral or slightly-curved texts but still suffers from highly-curved. One deficiency of point-based methods is that they usually predict a fixed number of points, which is hard to balance the performance between simple perspective texts and text instances with complex shapes (Zhang et al., 2022). ESIR (Zhan & Lu, 2019) adopts a single polynomial that fits the text center line to rectify the image for recognition. However, they can only model horizontal texts and still need points to represent text contours. In contrast, the polynomial curves in polynomial band can handle various orientations and shapes, from a straight line to a round curve with fewer and a fixed number of parameters.

Segmentation-based representation can naturally handle complex-shaped texts due to pixel-wise description (Wang et al., 2019a; Liao et al., 2021; Wang et al., 2019b; Tian et al., 2019; Zhu et al., 2021b; Liao et al., 2020; 2022). However, they frequently fail to divide adjacent texts due to ambiguous spatial positions. Although CentripetalText (Sheng et al., 2021) tackles this problem by detecting a shrunk text mask and reconstructing the contour by shift map, it suffers from high computation complexity. The proposed polynomial band overcomes this problem by considering the global curve shape. Curve coefficients in parameter space can be easily separated even though two texts are close, making more clear bounds in the text crowding scenario.

Text detection transformer. There is a trend to equip the transformer (Vaswani et al., 2017) with scene text detection. Current methods (Tang et al., 2022; Zhang et al., 2022) directly combine DETR variants (Carion et al., 2020; Zhu et al., 2021a) with point representation such as 16-point polygons or Bezier control points, and adopt the two-stage architecture for the ease of optimization. For example, FewBetter (Tang et al., 2022) first extracts segmentation maps by CNN-FPN (Lin et al., 2017) to show representative text regions, then samples feature points in each region and feed them into a transformer to further decode control points. TESTR (Zhang et al., 2022) follows (Wang et al., 2020b;c) to detect the bounding boxes, then utilize a transformer to find the control points in each bounding box. They sacrifice efficiency due to the two-stage pipelines and destroy the simplicity of the detection transformer scheme. Our PBFormer inherits the single-stage simplicity and inserts a parameter-free cross-scale pixel attention module between the CNN feature and transformer encoder-decoder layers. With the attention module, per-pixel features are fed into the subsequent transformer without generating any proposals, making the whole architecture more efficient.

3 METHODOLOGY

The overall framework of PBFormer is illustrated in Fig. 2. Given an image with texts, PBFormer first employs a ResNet50 to produce the multi-scale feature maps, then feed the feature maps to a cross-scale pixel attention module to highlight the texts' context information. The enhanced feature maps are concatenated and fed to a lightweight transformer, predicting the PBs' parameters. PB utilizes four polynomial curves to represent the shape of the text instance. It is simple but effective to capture different forms of texts. In the training phase, we sample the dense points from the curves for PB parameters estimation according to both predicted curve coefficients and the domain variables. A shape-constrained loss is designed to supervise the curves piece by piece.

3.1 NETWORK ARCHITECTURE.

The network contains three modules: a ResNet-based CNN encoder, a cross-scale pixel attention module, and a lightweight deformable transformer. We now introduce more details about the attention module and the decoders in the transformer.

Cross-scale pixel attention. The motivation of cross-scale pixel attention is to highlight text features in enlarged feature maps. To suppress the irrelevance from the other enlarged areas, it computes the attention weights of multi-scale features across both the scale and spatial dimensions,

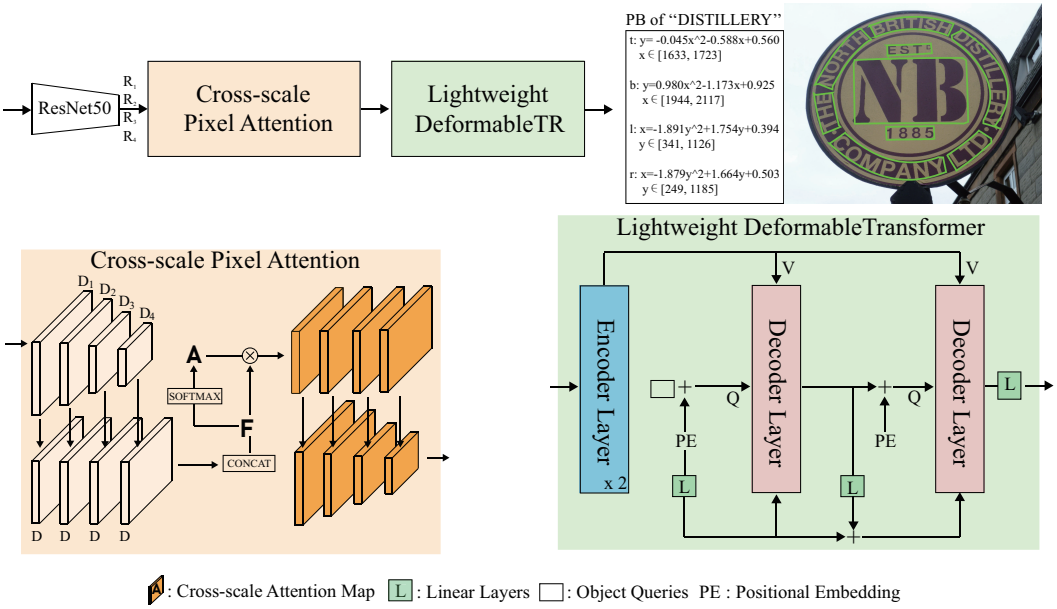

Figure 2: **Architecture of PBFormer.** The $\otimes$ and $+$ are element-wisely multiplication and addition, respectively. In the lightweight Deformable-Transformer, the object queries (white box) and Positional Embeddings (PE) are all learned parameters.

which implicitly down-weight the useless content. More details are illustrated in Fig. 2, the feature maps of a square image from the ResNet backbone are with the size R_1, R_2, R_3, R_4. We enlarge them to square feature maps with sizes D_1, D_2, D_3, D_4. Following Deformable-DETR, the four feature maps are transformed to have the same channels by 1×1 convolutions. Then we re-scaled them to have the same size D and assemble them to obtain $\mathbf{F} \in \mathbb{R}^{D \times D \times C \times 4}$. After that, we use a softmax layer to compute the attention map for each pixel and each channel: $\mathbf{A}^{ijk} = \text{softmax}([\mathbf{F}^{ijk1}, \mathbf{F}^{ijk2}, \mathbf{F}^{ijk3}, \mathbf{F}^{ijk4}]), \quad i \in [1, D], \quad j \in [1, D], \quad k \in [1, C]$. The attention map $\mathbf{A}$ and the feature map $\mathbf{F}$ have the same shape, *i.e.*, $\mathbf{A}, \mathbf{F} \in \mathbb{R}^{D \times D \times C \times 4}$. They multiply together to obtain the enhanced feature map $\mathbf{F}'$. We disassemble $\mathbf{F}'$ to four feature maps. All of them are with the size D and the channel dimension C. We re-scale their size to be D_1, D_2, D_3 and D_4. It is noteworthy that the whole cross-pixel attention module is parameter-free, which brings no training burdens during gradient back-propagation.

Lightweight deformable transformer. The four feature maps from the cross-scale pixel attention module are flattened to be vectors. We concatenate them to a long vector, then feed the vector to the deformable transformer. To predict the parameters of PB, we reduce the layers of the standard transformer encoder and decoder from 6 to 2, which is sufficient to yield competitive results. In a deformable transformer, the reference points attend a small set of key sampling points nearby for each query, which are important to the deformable attention module. We adopt a coarse-to-fine strategy to generate the reference points for the two decoder layers. In the first decoder, we adopt rough 2-d reference points derived from the positional embedding via a linear projection. In the second decoder, we combine the same 2-d reference points with a 2-d vectors transformed from the output of the first decoder. In particular, the 2-d vectors encode the relative offsets according to the first decoder's learned non-local dependencies, which help to generate more reasonable reference points for the second decoder layer. After that, a 3-layer MLP generates the PB predictions over the entire image.

3.2 Polynomial Band

We utilize four polynomial curves to represent the text instance's top, bottom, left, and right sides. The top and bottom boundaries are represented by $y = f^t(x)$ and $y = f^b(x)$:

$$\begin{aligned} y = f^t(x) &= a_2^t x^2 + a_1^t x + a_0^t, \quad x \in \left[e_0^t, e_1^t\right], \\ y = f^b(x) &= a_2^b x^2 + a_1^b x + a_0^b, \quad x \in \left[e_0^b, e_1^b\right], \end{aligned} \tag{1}$$

where (x, y) is the coordinate of a point on the boundary. $a_2^t, a_1^t, a_0^t, a_2^b, a_1^b, a_0^b$ are polynomial coefficients. $[e_0^t, e_1^t]$ and $[e_0^b, e_1^b]$ are range of x variable.

One critical problem is that the polynomial curve is a *single-value function* which means one point in the definition domain has a unique value in the value domain. The functions of the curves along the horizontal direction cannot be used to represent the curves along the vertical direction. For example, as Fig. 2 shows, the text 'DISTILLERT's left (or right) side would not be represented by any $y = f(x)$. We utilize $x = f^l(y)$ and $x = f^r(y)$ to represent the left and right polynomial curves of the text instance:

$$\begin{aligned} x = f^l(y) = a_2^l y^2 + a_1^l y + a_0^l, \quad y \in \left[e_0^l, e_1^l\right], \\ x = f^r(y) = a_2^r y^2 + a_1^r y + a_0^r, \quad y \in \left[e_0^r, e_1^r\right], \end{aligned} \tag{2}$$

where $[e_0^l, e_1^l]$ and $[e_0^r, e_1^r]$ define the range of y variable.

Output definition. We use four polynomial curves $y = f^t(x), y = f^b(x), x = f^l(y), x = f^r(y)$ that denote a band to wrap a text instance. Thus, the output is a 20-tuple θ that consists of all polynomial coefficients and boundary variables in the form of:

$$\theta = \left(a_2^t, a_1^t, a_0^t, e_0^t, e_1^t, a_2^b, a_1^b, a_0^b, e_0^b, e_1^b, a_2^l, a_1^l, a_0^l, e_0^l, e_1^l, a_2^r, a_1^r, a_0^r, e_0^r, e_1^r\right), \tag{3}$$

where $a_2^t, a_2^b, a_2^l, a_2^r \neq 0$, $e_0^t, ... \in [0, 1]$, and all of them are real numbers.

3.3 Loss Function

We propose shape-constrained loss to supervise the whole network. The network outputs N different PB parameters for each image, while their correspondences to ground truth contours are unknown. In this section, we first introduce how to compute the similarity between the predicted PB and ground truth contour by shape-constrained loss, then provide the loss function for the whole image based on optimized correspondences solved by bipartite matching.

The shape constraints for each curve. We first revisit the curve fitting loss without constraints used in the lane detection (Liu et al., 2021; 2022). The ground truth fitting points of a top or bottom curve are given by:

$$\hat{\mathcal{P}} = \{(\hat{x}_i, \hat{y}_i)\}_{i=0}^{K}, \quad \hat{x}_i = \hat{x}_0 + \frac{\hat{x}_K - \hat{x}_0}{K} i, \tag{4}$$

where the points are ordered from one end to the other, and the adjacent points for the top and bottom curves have the equal distance. The conventional fitting loss is:

$$\mathcal{L}_{w/o}(\hat{\mathcal{P}}) = \sum_{i=0}^{K} \|\hat{y}_i - f(\hat{x}_i)\|_1 + \|e_0 - \hat{x}_0\|_1 + \|e_1 - \hat{x}_K\|_1. \tag{5}$$

The fitting loss for the left and right curves can be obtained by exchanging the x and y variables. However, such a fitting loss is unsuitable for detecting texts with diverse shapes and different positions. It has two limitations: (1) the predicted curve segment is not aligned with the ground truth fitting points piece-by-piece; (2) the shape and range of the curve are independently optimized. As demonstrated in Fig. 3, the conventional loss is not sensitive to the length of the curves, therefore the text detector tends to detect curves with inaccurate lengths.

We consider to impose the shape constraints. The points on the predicted curve are sampled according to both curve shape and range:

$$\mathcal{P} = \{(x_i, f(x_i))\}_{i=0}^{K}, \quad x_i = e_0 + \frac{e_1 - e_0}{K} i. \tag{6}$$

Then we compare the predicted points $\mathcal{P}$ and ground truth points $\hat{\mathcal{P}}$:

$$\mathcal{L}(\mathcal{P}, \hat{\mathcal{P}}) = \sum_{i=0}^{K} \|x_i - \hat{x}_i\|_1 + \|f(x_i) - \hat{y}_i\|_1. \tag{7}$$

Leveraging Eq. 7 in text detection encourages PB to reconstruct the correct length of the contours.

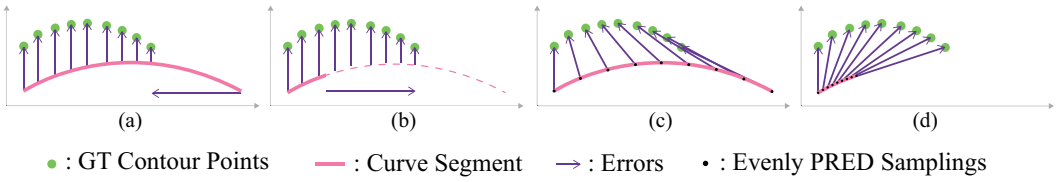

Figure 3: **Diagram of shape constraints.** All curve segments have the same curve coefficients. (a) and (c) have same ranges, so do (b) and (d). Without shape constraints, (a) and (b) show how to compare predicted curve segment with ground truth contour points. (c) and (d) illustrate the way with shape constraints.

The bipartite matching for the whole image. Let the network output of one image be $\mathcal{H} = \{h_j = (c_j, \theta_j)\}_{j=1}^{N}$, where c_j is the confidence score indicating the possibility of a PB covering a text and N is set to be larger than the maximum number of texts in an image. After sampling the points on the four curves according to Eq. 6, $\mathcal{H}$ can be further represented by: $\mathcal{H} = \{h_j = (c_j, \mathcal{P}_j^t, \mathcal{P}_j^b, \mathcal{P}_j^l, \mathcal{P}_j^r)\}_{j=1}^{N}$.

For bipartite matching, we pad the ground truth set $\hat{\mathcal{H}}$ with non-text instances to have a size N. The element having text instance is represented by $\hat{h}_j = (\hat{c}_j, \hat{\mathcal{P}}_j^t, \hat{\mathcal{P}}_j^b, \hat{\mathcal{P}}_j^l, \hat{\mathcal{P}}_j^r)$. In particular, $\hat{\mathcal{P}}_j^t, \hat{\mathcal{P}}_j^b, \hat{\mathcal{P}}_j^l, \hat{\mathcal{P}}_j^r$ are sampled points according to Eq. 4, while they are not need to be instantiated in the matching cost for non-text instances thus are set to be $\emptyset$. $\hat{c}_j$ is set to be 1 for the text and 0 for the non-text class. We formulate a bipartite matching problem to find an optimal injective function $g : \hat{\mathcal{H}} \rightarrow \mathcal{H}$, *i.e.*, $g(i)$ is the index of the PB assigned to fitting the i-th ground truth text:

$$g^* = \arg\min_{g} \sum_{j=1}^{N} \mathcal{L}^{fit}\left(\hat{h}_j, h_{g(j)}\right) + \mathcal{C}^{focal}\left(\hat{c}_j, c_{g(j)}\right), \tag{8}$$

where $\mathcal{L}^{fit}$ is the fitting loss and $\mathcal{C}^{focal}$ is the focal cost. The fitting loss compares the predicted contour and ground truth contour by using the loss defined in Eq. 7:

$$\mathcal{L}^{fit}\left(\hat{h}_j, h_{g(j)}\right) = \mathbf{1}_{\hat{c}_j>0}\left(\mathcal{L}(\hat{\mathcal{P}}_j^t, \mathcal{P}_{g(j)}^t) + \mathcal{L}(\hat{\mathcal{P}}_j^b, \mathcal{P}_{g(j)}^b) + \mathcal{L}(\hat{\mathcal{P}}_j^l, \mathcal{P}_{g(j)}^l) + \mathcal{L}(\hat{\mathcal{P}}_j^r, \mathcal{P}_{g(j)}^r)\right). \tag{9}$$

Then, the focal cost is defined as the difference between the positive and negative costs:

$$\mathcal{C}^{focal}\left(\hat{c}_j, c_{g(j)}\right) = \lambda \mathbf{1}_{\hat{c}_j>0}\left[-\alpha\left(1 - c_{g(j)}\right)^{\gamma} \log c_{g(j)} + (1-\alpha)\, c_{g(j)}^{\gamma} \log\left(1 - c_{g(j)}\right)\right], \tag{10}$$

where α and γ are the hyper-parameter for the focal loss. α is used to address the class imbalance, and γ adjusts the rate at which easy examples are down-weighted. λ adjusts the weight of the focal cost. The bipartite problem (Eq. 8) can be efficiently solved by the Hungarian algorithm.

Overall Loss. With the optimized g^*, the overall loss function is given by:

$$\mathcal{L}^{overall} = \sum_{j=1}^{N} \mathcal{L}^{fit}\left(\hat{h}_j, h_{g^*(j)}\right) + \mathcal{L}^{focal}\left(\hat{c}_j, c_{g^*(j)}\right), \tag{11}$$

where the $\mathcal{L}^{focal}\left(\hat{c}_j, c_{g^*(j)}\right)$ is the focal loss:

$$\mathcal{L}^{focal}\left(\hat{c}_j, c_{g^*(j)}\right) = \lambda\left[\mathbf{1}_{\hat{c}_j>0} - \alpha\left(1 - c_{g^*(j)}\right)^{\gamma} \log c_{g^*(j)} - \mathbf{1}_{\hat{c}_j=0}\,(1-\alpha)\, c_{g^*(j)}^{\gamma} \log\left(1 - c_{g^*(j)}\right)\right]. \tag{12}$$

α, λ and γ are the same with the ones in Eq. 10.

4 Experiments

Datasets. CTW1500 (Liu et al., 2019) is a multi-oriented and curved scene text detection benchmark containing 1,000 training and 500 testing images. Annotations are based on the text-line level with fixed fourteen points. The majority of text instances are curved. Total-Text (Chng & Chan, 2017) is an another multi-oriented and curved scene text benchmark, while it consists of various text shapes such as multidirectional quadrilateral. It has 1255 training images and 300 testing images. Each instance is annotated by ten point text-line.

Table 1: Detection results on CTW1500 and Total-Text **without** pre-training on any text datasets. "Rep." denotes the method's output representation. "F.", "Prec.", "Rec." represent F-measure, Precision, and Recall. All the results are from their official codes and models. TESTR* means training without recognition branch.

		CTW1500				Total-Text			
Method	Rep.	F.	Prec.	Rec.	FPS	F.	Prec.	Rec.	FPS
PSENet	Seg	78.0	80.6	75.6	3.9	78.3	81.8	75.1	3.9
PAN	Seg	81.0	84.6	77.7	39.8	83.5	88.0	79.4	39.6
FCENet	Seg	85.1	88.1	82.3	2.7	85.8	89.3	82.5	2.9
ContourNet	Pts	83.9	83.7	84.1	3.8	85.4	86.9	**83.9**	3.8
TextBPN	Pts	84.0	87.7	80.6	12.1	86.9	90.8	83.3	10.6
TESTR*	Pts	85.1	88.4	82.1	5.6	85.3	89.7	81.2	5.3
TESTR*	Bez	84.7	87.9	81.8	5.6	86.3	90.3	82.6	5.5
PBFormer	PB	**87.0**	**89.6**	**84.5**	24.7	**87.1**	**92.1**	82.6	24.6

Table 2: Detection results on CTW1500 and Total-Text **with** pre-training on text datasets. MLT, ST, ArT, and CST are abbreviations for MLT2017, SynthText, ArT 2019 and CurvedSynthText datasets. C+M+T means using a combination of CST, MLT, and Total-Text for pre-training.

			CTW1500				Total-Text			
Method	Rep.	Ext.	F.	Prec.	Rec.	FPS	F.	Prec.	Rec.	FPS
1 PSENet	Seg	MLT	82.2	84.8	79.7	3.9	80.9	84.0	78.0	3.9
PAN	Seg	ST	83.7	86.4	81.2	39.8	85.0	89.3	81.0	39.6
DB	Seg	ST	83.4	86.9	80.2	22	84.7	87.1	82.5	32
DB++	Seg	ST	85.3	87.9	82.8	26	86.0	88.9	83.2	28
TextRay	Pts	ArT	81.6	82.8	80.4	3.2	80.6	83.5	77.9	3.5
DRRG	Pts	MLT	84.5	85.9	83.0	-	85.7	86.5	84.9	-
TextBPN	Pts	MLT	85.0	86.5	83.6	12.2	87.9	90.7	85.2	10.7
TESTR*	Pts	C+M+T	86.6	**90.8**	82.8	5.6	86.2	92.4	80.7	5.3
ABCNet	Bez	CST	81.4	84.4	78.5	6.8	84.5	87.9	81.3	6.9
FewBetter	Bez	CST	85.2	88.1	82.4	-	**88.1**	90.7	**85.7**	-
TESTR*	Bez	C+M+T	85.9	90.6	81.6	5.6	87.4	92.4	82.8	5.5
PBFormer	PB	CST	**88.0**	90.6	**85.4**	24.7	**88.1**	**93.2**	83.5	24.6

Evaluation Metrics. We follow the standard metrics F-measures, recall, and precision to evaluate the performance. A prediction is considered as a true positive only when its IoU from the corresponding ground truth contour is larger than 0.5.

Implementation Details. The input image size is set to be 800×800 for training and testing. Loss coefficient α, γ and λ are set as 0.25, 2 and 2. The fixed number of output N is 300. In the cross-scale pixel attention module, R_1, R_2, R_3, R_4 are 100, 50, 25, 13, D_1, D_2, D_3, D_4 are set as 128, 64, 32, 16, and we set $D = D_2$. For training from scratch, the learning rate is set to be 1×10^{-4} and decayed ten times at 7200 epochs, and the total number of training iterations is set as 9000 epochs. The training process takes about 2 days on 4 Tesla V100 GPUs with the image batch size of 14. For training with pre-training, we pre-train the model for 50 epochs, then fine-tune the model on CTW1500 and Total-Text by the same setting as training without pre-training states.

4.1 Comparison with the State-of-the-art Methods

To demonstrate the effectiveness of our method, we take (1) point-based methods ABCNet (Liu et al., 2020), TextRay (Wang et al., 2020a), DRRG (Zhang et al., 2020), ContourNet (Wang et al., 2020c) and TextBPN (Zhang et al., 2021); (2) segmentation-based methods PSENet (Wang et al., 2019a), PAN (Wang et al., 2019b), DB (Liao et al., 2020), FCENet (Zhu et al., 2021b), and DB++ (Liao et al., 2022); and (3) recent transformer-based methods FewBetter (Tang et al., 2022) and TESTR (Zhang et al., 2022). For a fair comparison with TESTR, we use their official training codes, settings, and models but only set the recognition loss to zeros. We will not use text character annotation for text detection training.

Figure 4: **Qualitative comparisons with previous SOTA on Total-Text and CTW1500.** Compared to DB++ and FCENet, our PBFormer predicts more compact and precise contours for crowded texts (the first two are DB++'s Total-Text detections, and the last four are FCENet's CTW1500 detections, because they did not release the model of another dataset). Compared to TESTR, PBFormer reduces false negatives and performs better for long and curved texts.

Figure 5: **Effect visualization of shape-constrained loss**. For each image pair, the left image shows the results with the fitting loss Eq. 7, and the right image is with the shape-constrained loss Eq. 11. With the shape-constrained loss, PBFormer outputs more complete contours.

Figure 6: **Visualization of CPA's attention maps.**The left four images show CPA produces attention for the small texts at a swallow layer and the large texts at a deep layer. The right four image show attention is concentrated on one layer due to texts having similar sizes.

Models trained from scratch. As shown in Tab. 1, PBFormer establishes a new state-of-the-art of **87.0**% on CTW1500, which is **1.9**% better than previous best TESTR while achieving **4.4** $\times$ FPS. Moreover, PBFormer yields the best F-measure **87.1**% on Total-Text, which is **0.2**% better than previous best TextBPN while being **2.3** $\times$ FPS. Compared to TESTR, PBFormer also improves it by **0.8**% while keeping a **4.5** $\times$ FPS.

Models with pre-training. Since previous methods have different choices of pre-training datasets, we choose to use the CurvedSynthText as FewBetter did. As shown in Tab. 2, PBFormer can achieve better results from the model pre-training. On the CTW1500 dataset, our method achieves the best results. It outperforms previous best TESTR by **1.4**% in terms of F-measure and is **4.4** $\times$ faster than TESTR. On the Total-Text dataset, PBFormer also has the best performance. Compared to TESTR, PBFormer performs a **4.5** $\times$ FPS while being **0.7**% higher F-measure. In addition, PBFormer yields **2.7**% and **2.1**% F-measure better than the previous best segmentation-based DB++ on CTW1500 and Total-Text, while keeping a very close FPS performance.

Qualitative comparisons. Considering crowded texts in Fig. 4(a),(e), and (f), PBFormer performs fewer false-negatives than TESTR and more accurate contours than DB++ and FCENet. When texts have very long shapes or have characters' large scale-changes, PBFormer detected more completed contours than DB++, FCENet, and TESTR, as Fig. 4(d) and (c) have shown.

Table 3: CPA's influences and comparisons with FPN and ASF. Ext. means trainable parameters.

Module	Ext.	F.	Prec.	Rec.
-	-	86.0	90.5	81.9
Enlarge	-	85.2	88.9	81.8
Attn.	-	86.1	90.6	82.1
CPA	-	**87.1**	**92.1**	**82.6**
ASF	✓	85.3	90.9	80.4
FPN	✓	85.0	89.5	81.0

Table 4: Comparisons of the number of different supervision points per polynomial curve.

Number	F.	Prec.	Rec.
6	84.6	90.5	79.5
12	86.0	90.4	81.9
24	**87.1**	**92.1**	**82.6**
30	86.5	91.1	82.4
36	86.4	91.1	82.3

4.2 Ablation Study

Effect of cross-scale pixel attention(CPA). We first analyze the effect of the CPA module. Tab. 3 shows the performance of the models with different CPA configurations and the models replacing CPA with other fusion modules. The study of the CPA configuration has the following conclusions: (1) just enlarging features performs 0.8% worse; (2) only using attention has a minor 0.1% improvement; (3) combining both boosts the performance by 1.1% significantly. The effectiveness of the CPA module is that the attentional fusion adaptively highlights texts' features at a suitable scale and suppresses the features of other scales. Fig. 6 demonstrates the four attention maps across scales from shallow to deep layers of the backbone. Moreover, CPA performs 1.8% and 2.1% better than FPN and ASF (Liao et al., 2022). We attribute it to the FPN being too heavy to train well, and ASF might distort features by additional convolutions.

Effect of shape-constrained loss. We now analyze the influence of shape-constrained loss. In Tab. 5, we can observe that model with shape constraints can improve the model without the constraints by a significant **3.6**% in terms of F-measure. As Fig. 5 shows, the model trained with shape-constrained loss can produce more complete contours (the right image in the pair) than the models without the loss (the left image in the pair).

Table 5: Effect of shape constrained loss.

Constraints	F.	Prec.	Rec.
-	83.5	87.5	79.9
✓	**87.1**	**92.1**	**82.6**

Table 6: Comparisons of polynomial's order.

Curve	F.	Prec.	Rec.
Quadratic	**87.1**	92.1	82.6
Cubic	87.0	**92.2**	82.4
Quartic	86.7	91.0	**82.7**

Investigation of PB's configurations. We first analyze the influence of the polynomial order of PB. Theoretically, higher-order polynomials could fit more complex texts, but the quadric curves can describe the texts in most current datasets well. For this reason, the higher-order curves are overqualified to represent the text shapes in the existing datasets, but they are easily overfitted. Secondly, we analyze the effect of sampling points' density. As Tab. 4 shows, too few points perform 2.5% worse since the quantity is insufficient to learn curved shapes, while too many points also decrease 0.7% due to redundant learning points, especially for straight texts.

5 Conclusion

We have presented PBFormer, an efficient and accurate text detection method. It is superior to handle crowded texts or texts with diverse shapes. PBFormer equips a new text representation, Polynomial Band, to a transformer-based network consisting of a cross-scale pixel attention module and a lightweight deformable transformer. We supervise the network with a shape-constrained loss term, encouraging the network to output the correct contour length. PBFormer shows strong robustness when training without pre-training on the additional datasets, which is much more resource-friendly than other transformer-based methods.

REFERENCES

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pp. 213–229, 2020.

Chee Kheng Chng and Chee Seng Chan. Total-text: A comprehensive dataset for scene text detection and recognition. In *ICDAR*, pp. 935–942. IEEE, 2017.

Hui Li, Peng Wang, and Chunhua Shen. Towards end-to-end text spotting with convolutional recurrent neural networks. In *ICCV*, pp. 5248–5256. IEEE Computer Society, 2017.

Hui Li, Peng Wang, and Chunhua Shen. Towards end-to-end text spotting in natural scenes. *CoRR*, abs/1906.06013, 2019.

Minghui Liao, Zhaoyi Wan, Cong Yao, Kai Chen, and Xiang Bai. Real-time scene text detection with differentiable binarization. In *AAAI*, pp. 11474–11481. AAAI Press, 2020.

Minghui Liao, Pengyuan Lyu, Minghang He, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 43(2):532–548, 2021.

Minghui Liao, Zhisheng Zou, Zhaoyi Wan, Cong Yao, and Xiang Bai. Real-time scene text detection with differentiable binarization and adaptive scale fusion. *CoRR*, abs/2202.10304, 2022.

Tsung-Yi Lin, Piotr Dollár, Ross B. Girshick, Kaiming He, Bharath Hariharan, and Serge J. Belongie. Feature pyramid networks for object detection. In *CVPR*, pp. 936–944. IEEE Computer Society, 2017.

Ruijin Liu, Zejian Yuan, Tie Liu, and Zhiliang Xiong. End-to-end lane shape prediction with transformers. In *WACV*, pp. 3693–3701, 2021.

Ruijin Liu, Dapeng Chen, Tie Liu, Zhiliang Xiong, and Zejian Yuan. Learning to predict 3d lane shape and camera pose from a single image via geometry constraints. In *AAAI*, pp. 1765–1772. AAAI Press, 2022.

Yuliang Liu, Lianwen Jin, Shuaitao Zhang, Canjie Luo, and Sheng Zhang. Curved scene text detection via transverse and longitudinal sequence connection. *Pattern Recognit.*, 90:337–345, 2019.

Yuliang Liu, Hao Chen, Chunhua Shen, Tong He, Lianwen Jin, and Liangwei Wang. Abcnet: Real-time scene text spotting with adaptive bezier-curve network. In *CVPR*, pp. 9806–9815. Computer Vision Foundation / IEEE, 2020.

Shangbang Long, Xin He, and Cong Yao. Scene text detection and recognition: The deep learning era. *Int. J. Comput. Vis.*, 129(1):161–184, 2021.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8748–8763. PMLR, 2021.

Sangeeth Reddy, Minesh Mathew, Lluís Gómez, Marçal Rusiñol, Dimosthenis Karatzas, and C. V. Jawahar. Roadtext-1k: Text detection & recognition dataset for driving videos. In *ICRA*, pp. 11074–11080. IEEE, 2020.

Tao Sheng, Jie Chen, and Zhouhui Lian. Centripetaltext: An efficient text instance representation for scene text detection. In *NeurIPS*, pp. 335–346, 2021.

Yipeng Sun, Chengquan Zhang, Zuming Huang, Jiaming Liu, Junyu Han, and Errui Ding. Textnet: Irregular text reading from images with an end-to-end trainable network. In *ACCV (3)*, pp. 83–99. Springer, 2018.

Jingqun Tang, Wenqing Zhang, Hongye Liu, Mingkun Yang, Bo Jiang, Guanglong Hu, and Xiang Bai. Few could be better than all: Feature sampling and grouping for scene text detection. In *CVPR*. Computer Vision Foundation / IEEE, 2022.

Zhuotao Tian, Michelle Shu, Pengyuan Lyu, Ruiyu Li, Chao Zhou, Xiaoyong Shen, and Jiaya Jia. Learning shape-aware embedding for scene text detection. In *CVPR*, pp. 4234–4243. Computer Vision Foundation / IEEE, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pp. 5998–6008, 2017.

Fangfang Wang, Yifeng Chen, Fei Wu, and Xi Li. Textray: Contour-based geometric modeling for arbitrary-shaped scene text detection. In *ACM Multimedia*, pp. 111–119. ACM, 2020a.

Hao Wang, Pu Lu, Hui Zhang, Mingkun Yang, Xiang Bai, Yongchao Xu, Mengchao He, Yongpan Wang, and Wenyu Liu. All you need is boundary: Toward arbitrary-shaped text spotting. In *AAAI*, pp. 12160–12167. AAAI Press, 2020b.

Wenhai Wang, Enze Xie, Xiang Li, Wenbo Hou, Tong Lu, Gang Yu, and Shuai Shao. Shape robust text detection with progressive scale expansion network. In *CVPR*, pp. 9336–9345. Computer Vision Foundation / IEEE, 2019a.

Wenhai Wang, Enze Xie, Xiaoge Song, Yuhang Zang, Wenjia Wang, Tong Lu, Gang Yu, and Chunhua Shen. Efficient and accurate arbitrary-shaped text detection with pixel aggregation network. In *ICCV*, pp. 8439–8448. IEEE, 2019b.

Yuxin Wang, Hongtao Xie, Zheng-Jun Zha, Mengting Xing, Zilong Fu, and Yongdong Zhang. Contournet: Taking a further step toward accurate arbitrary-shaped scene text detection. In *CVPR*, pp. 11750–11759. Computer Vision Foundation / IEEE, 2020c.

Fangneng Zhan and Shijian Lu. ESIR: end-to-end scene text recognition via iterative image rectification. In *CVPR*, pp. 2059–2068. Computer Vision Foundation / IEEE, 2019.

Shi-Xue Zhang, Xiaobin Zhu, Jie-Bo Hou, Chang Liu, Chun Yang, Hongfa Wang, and Xu-Cheng Yin. Deep relational reasoning graph network for arbitrary shape text detection. In *CVPR*, pp. 9696–9705. Computer Vision Foundation / IEEE, 2020.

Shi-Xue Zhang, Xiaobin Zhu, Chun Yang, Hongfa Wang, and Xu-Cheng Yin. Adaptive boundary proposal network for arbitrary shape text detection. In *ICCV*, pp. 1285–1294. IEEE, 2021.

Xiang Zhang, Yongwen Su, Subarna Tripathi, and Zhuowen Tu. Text spotting transformers. In *CVPR*. Computer Vision Foundation / IEEE, 2022.

Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: deformable transformers for end-to-end object detection. In *ICLR*. OpenReview.net, 2021a.

Yiqin Zhu, Jianyong Chen, Lingyu Liang, Zhanghui Kuang, Lianwen Jin, and Wayne Zhang. Fourier contour embedding for arbitrary-shaped text detection. In *CVPR*, pp. 3123–3131. Computer Vision Foundation / IEEE, 2021b.