

Progressive Evolution from Single-Point to Polygon for Scene Text

Linger Deng^{1†} Mingxin Huang^{2†} Xudong Xie¹ Yuliang Liu^{1*} Lianwen Jin² Xiang Bai¹
¹Huazhong University of Science and Technology ²South China University of Technology

lingerdeng2023@163.com ylliu@hust.edu.cn

Abstract

The advancement of text shape representations towards compactness has enhanced text detection and spotting performance, but at a high annotation cost. Current models use single-point annotations to reduce costs, yet they lack sufficient localization information for downstream applications. To overcome this limitation, we introduce Point2Polygon, which can efficiently transform single-points into compact polygons. Our method uses a coarse-to-fine process, starting with creating and selecting anchor points based on recognition confidence, then vertically and horizontally refining the polygon using recognition information to optimize its shape. We demonstrate the accuracy of the generated polygons through extensive experiments: 1) By creating polygons from ground truth points, we achieved an accuracy of 82.0% on ICDAR 2015; 2) In training detectors with polygons generated by our method, we attained 86% of the accuracy relative to training with ground truth (GT); 3) Additionally, the proposed Point2Polygon can be seamlessly integrated to empower single-point spotters to generate polygons. This integration led to an impressive 82.5% accuracy for the generated polygons. It is worth mentioning that our method relies solely on synthetic recognition information, eliminating the need for any manual annotation beyond single points. The code will be publicly available.

1. Introduction

In recent years, the evolution of text shape representation has significantly enhanced text detection and spotting, progressing from managing horizontal and multi-oriented text to adeptly handling arbitrarily shaped text. However, polygonal annotation brings expensive annotation costs and limits the acquisition of large-scale annotated data, constraining the overall generalization of the model.

To reduce the requirement of costly annotation, some researchers have delved into weakly supervised methods.

[†]Equal contribution.

^{*}Corresponding author.




Figure 1. Point2Polygon maintains cost-efficient point annotation while automatically generating polygons with high accuracy.

Several approaches [1, 8, 30, 36, 39] employ coarse detection annotations to yield refined detection outcomes. Another strategy [14] combines synthetic data with box annotations and real data with recognition annotations. Recent research by Liu *et al.* [24] demonstrated that employing cost-effective single-point annotations for location supervision can yield competitive results in text spotting. Nonetheless, this approach has limitations in certain downstream applications. Kil *et al.* [13] argue that single-point annotations fall short in providing adequate text location information for tasks like scene text editing [27, 35] and text removal [19, 34].

Hence, given the high cost of polygonal annotation and the necessity for accurate text representation, a key question arises, as shown in Fig. 1: *Can we accurately generate text polygon representations using only single-point annotation, to keep both the advantages of cost-effective annotation and compact representation?*

However, deriving an accurate polygon directly from weak supervision of single-point could be a significant chal-



Figure 2. Adapting to the development in text shape representation, we present Point2Polygon, a novel method that significantly alleviates the constraints of point supervision.

lenge due to the considerable variations in shape, size, style, distortion, and layout inherent in scene text. However, recent studies [3, 9, 10, 28, 37] suggest recognition information can be beneficial for predicting accurate detection results.

In this paper, we present Point2Polygon, an efficient method to transform single points into polygons using a three-step, coarse-to-fine process. Initially, our Anchor Generation Module (AGM) creates multiple anchors from input image sizes and selects the best one using initial recognition confidence. Next, the Polygon Generation Module (PGM) shapes these anchors into vertical boundary points, guided by recognition loss and supervised by the Thin-Plate Spline technique. The final step involves our Polygon Rectification Module (PRM), which adjusts the polygon horizontally, using detailed recognition information to remove excess points beyond the region of the text. Extensive experiments demonstrate the accuracy of the generated polygons. We summarize the advantages of the proposed method as follows:

- We introduce a novel framework, termed Point2Polygon, which can efficiently leverage multi-granularity recognition information to evolve points into polygons.
- Point2Polygon requires only single-point annotations and synthetic recognition data. Its simplicity also allows it to be seamlessly integrated with single-point spotters like SPTS v2 for autonomous polygon generation, with an impressive 82.5% accuracy.
- Experiments demonstrate the accuracy of the generated polygon through multiple aspects: It achieves 82.0% accuracy on ICDAR 2015 by generating polygons from ground truth points, and an average of 86% accuracy training using several detectors.

2. Related Work

In recent years, text representation of scene text has progressed through distinct stages of development, as presented in Fig. 2. Early methods mainly focused on horizontal text detection. For instance, Tian *et al.* [31] introduced a CTPN that predicts a sequence of fixed-width text fragments along with a vertical anchor frame. Liao *et*

al. [16] accurately adjust the anchor box and classification layer based on the SSD [21].

Subsequently, the research focus has shifted towards multi-oriented text detection. For example, Zhou *et al.* [41] proposed EAST to regress the four points in multi-oriented text. He *et al.* [7] proposed a framework to directly regress the multi-oriented text.

Although significant progress has been achieved, researchers find the previous methods still fall short in handling arbitrarily shaped text. To address this issue, two categories have been developed. The first treats text detection as a segmentation task, initially segmenting the text kernel and then expanding to the full text. For example, Wang *et al.* [32] proposed PSENet, which progressively expands the kernel to the final text boundary. For real-time detection, Wang *et al.* [33] proposed a learnable post-processing method to expand the kernel. Liao *et al.* [18] introduced an adaptive binarization process for the segmentation model. In contrast to segmentation-based methods, regression-based methods utilize curves directly to fit text. Liu *et al.* [23] proposed a parameterized Bezier curve to fit the arbitrarily shaped text. Similarly, Liu *et al.* [20] adopt polynomials to represent the curve text.

Although arbitrarily shaped text can precisely represent the shape of the text, the cost of annotation is expensive. Therefore, some researchers have begun to explore reducing the cost of annotation. One solution utilizes the coarse detection annotation to generate fine detection results. Hu *et al.* [8] and Tian *et al.* [30] have proposed a similar solution, employing word or line-level annotation to achieve character-level detection. Additionally, Wu *et al.* [36] have utilized curved synthetic data to train a segmentation module tailored for text with arbitrary shapes. Another solution involves using synthetic data with box annotation and real data with recognition annotation. Kittenplon *et al.* [14] propose a weakly-supervised method that uses synthetic data with box annotations and real data with text annotations to train the text spotting model. Despite these advancements, it is important to note that both approaches still require the expensive annotation of bounding boxes.

Therefore, some methods [24, 26, 29] attempt to further reduce the annotation costs by employing a single point,

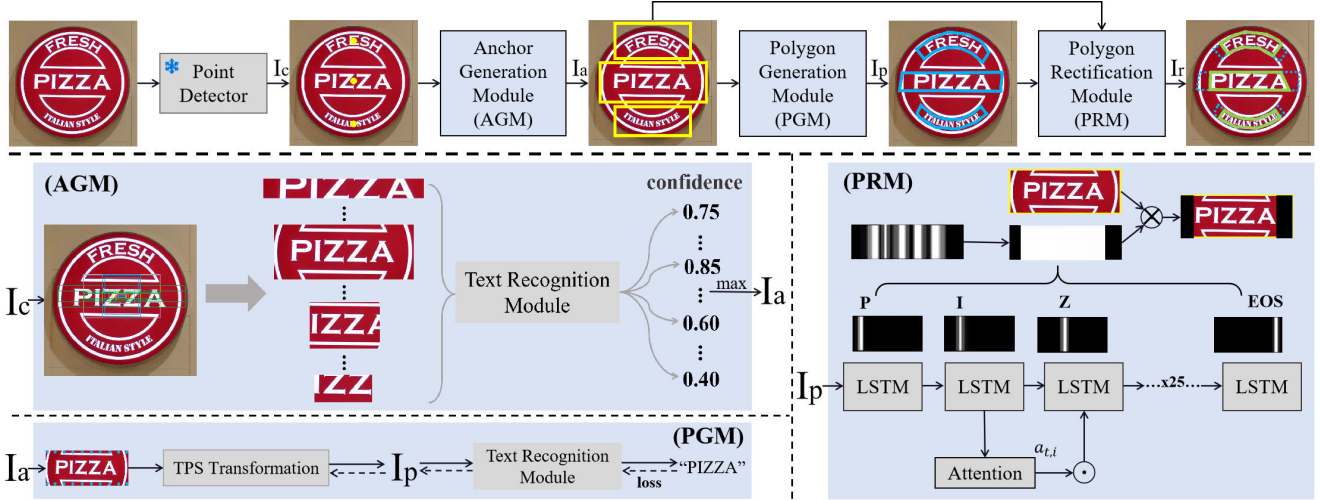


Figure 3. Overview of the proposed Point2Polygon model. We use the point detector as a basis for obtaining the final text polygon by supervising it from coarse to fine strategy.

which has shown significant success in end-to-end text spotting. However, it’s important to note that while a single point is effective in end-to-end text spotting, the bounding box remains crucial for numerous downstream tasks, as emphasized in [13]. In this paper, we introduce Point2Polygon, a novel approach that generates polygons using only single points and synthetic recognition information.

3. Methodology

In this paper, we present a solution for evolving polygons from a single point using coarse-to-fine recognition information. The overall pipeline is illustrated in Fig. 3. Firstly, given an image, a point detector is employed to locate the central point I_c of a text. Then, we introduce an Anchor Generation Module (AGM) to generate adaptive multiple anchors centered on these points and select the optimal one I_a . It is selected based on coarse-grained recognition information (recognition confidence). Next, these selected anchors are fed into the Polygon Generation Module (PGM) to refine the upper and lower boundary points I_p using the Thin-Plate spline [2]. This module operates under medium-grained recognition information (recognition loss). Finally, since the boundary points obtained in PGM often extend beyond the length of the text horizontally, we introduce the Polygon Rectification Module (PRM). PRM leverages fine-grained recognition information (attention map) to finely adjust the boundary points horizontally, producing a more precise adjustment I_r of the boundary points.

3.1. Anchor Generation

The Anchor Generation Module (AGM) intercepts multiple anchors centered on the text’s central point. Subsequently, these anchors are input into the text recognition module,

and the anchor with the highest recognition confidence is selected as the optimal anchor. The Fig. 3 illustrates the process of this module.

As the scene text exhibits significant variations in shape, we preset different default anchors to cover them. This includes “extra-long” default anchors designed for extra-large aspect ratio text, “large” anchors for large aspect ratio text, “normal” anchors for moderate aspect ratio text, and “short” anchors for small aspect ratio text. Each category is diversified into multiple sizes to effectively accommodate different text instances. Concretely, we empirically establish 4 sizes for extra-long default anchors, 6 sizes for long default anchors, 5 sizes for normal default anchors, and 6 sizes for short default anchors. Given the image size (w, h) , the sizes of the corresponding anchors are represented as follows:

$$(w_{sl}, h_{sl}) = \left(\frac{2w}{3} \text{ or } \frac{2w}{5}, \frac{h}{5q} \right), \quad (1)$$

$$(w_l, h_l) = \left(\frac{2w}{5j}, \frac{h}{5i} \right), \quad (2)$$

$$(w_n, h_n) = \left(\frac{2w}{5k}, \frac{2h}{5k} \right), \quad (3)$$

$$(w_s, h_s) = \left(\frac{w}{5i}, \frac{2h}{5j} \right), \quad (4)$$

where (w_{sl}, h_{sl}) , (w_l, h_l) , (w_n, h_n) , (w_s, h_s) are extra-long, long, normal, and short default anchors, respectively. i, j, k, q are scaling factors sets, where $q = 1, 2, 3, 4$, $i = 1, 2, \dots, 6$, $j = 1, 2, 4, 6, 8, 10$, and $k = 1, 2, 3, 6, 10$.

Then, we send the coarse anchors into the off-the-shell text recognition model and obtain the recognition confidence. The recognition confidence of an anchor is the sum of the confidences of each character in the recognition re-

sult: $P = \sum_i^n P_i/n$, where n is the number of characters contained in the text and P_i is the prediction confidence for each character. We select the optimal anchor with the highest recognition confidence.

3.2. Polygon Generation

To effectively represent arbitrarily shaped text, we propose a Polygon Generation Module (PGM) to refine the coarse anchors from AGM. The entire procedure is illustrated in the bottom left corner of Fig. 3. Firstly, we evenly sample 10 points along the upper and lower contour for each coarse anchor. Subsequently, we employ the Thin-Plate Spline (TPS) to map these sampled points to the boundaries of the text region. TPS transformation necessitates establishing a correspondence for each coordinate, enabling the transformation from the rectangle to the arbitrary text shape. We refer to the target shape of the text as \mathbf{G} and the predefined rectangle as \mathbf{O} . According to the TPS transform, the corresponding point of $p(x, y) \in \mathbf{G}$ on \mathbf{O} can be calculated by

$$\Phi(p) = m_0 + \mathbf{M} \cdot p + \sum_{i=1}^N \omega_i U(\|p - p_i\|), \quad (5)$$

where M and m_0 are unknown parameters to be calculated, N is the number of control points, $\|p - p_i\|$ is the distance from $p(x, y)$ to fixed points $p_i(x_i, y_i)$ on the target shape G . U is the radial basis function $U(r) = r^2 \ln r$, and w_i denotes the weighting for different radial bases. Then the TPS transform function is determined by the parameters

$$\mathbf{T} = [m_0 \quad M \quad w_1 \quad \dots \quad w_N], \quad (6)$$

whose shape is $[2, N + 3]$.

Through the utilization of the TPS transform function, the coarse anchors can be transformed into boundary points that are in closer proximity to the text region, enhancing the precision of localization. Inspired by the text recognition model [5, 25, 28, 40], we use recognition loss to supervise the learning of the locations of target boundary points on \mathbf{G} .

3.3. Polygon Rectification

We observed that the boundary points generated by the PGM are often excessively wide. To rectify this issue, we introduce a Polygon Rectification Module (PRM). The PRM leverages fine-grained recognition information (attention map), to remove the excess points beyond the region of the text. Following the TPS transformation, this problem does not manifest in the vertical direction. Therefore, we employ a one-dimensional attention map to refine the boundary points from the horizontal direction. The entire procedure is illustrated in the bottom right corner of Fig. 3. We crop the image using the boundary points obtained from PGM. Subsequently, the cropped regions are

forwarded to the recognizer, which generates a sequence of attention maps a_t highlighting the character regions in time steps t .

After obtaining the attention maps of each time step, we use a threshold τ to transform the attention maps. This process is formalized as:

$$\mathbf{a}_t = \begin{cases} 1, & \mathbf{a}_t > \tau, \\ 0, & \mathbf{a}_t < \tau. \end{cases} \quad (7)$$

We assign attention weights of 1 to steps where the attention weights surpass a specified threshold τ , and 0 to the rest. The attention map efficiently locates the character positions in the feature map. As a result, with the attention map, we can exclude points with zero attention weight, eliminating unnecessary information.

4. Experiments

To evaluate the effectiveness of our method, we conduct experiments on the popular scene text detection benchmarks, including ICDAR 2015 [12], TotalText [4], and SCUT-CTW1500 [22].

4.1. Implementation Details

In this paper, we specifically select ASTER [28] as our text recognition model. We utilize the synthetic datasets Synth90k [11] and SynthText [6] for training our text recognizer. We employ the Adadelta optimizer [38] and initially set the learning rate to 1.0, which is later reduced to 0.1 at step 0.6M and further to 0.01 at step 0.8M.

4.2. Datasets

ICDAR 2015. This dataset consists of 1000 training images and 500 test images of incidental scenes with complex backgrounds and multidirectional text. Some text can also appear in any orientation and at any location at small or low resolution, with high variability in blurring, distortion, etc.

TotalText. This dataset is designed to be used for arbitrarily-shaped text detection and recognition, which contains 1,255 training images and 300 test images. The images are annotated with word-level annotations.

SCUT-CTW1500. This dataset is an arbitrarily-shaped scene text benchmark in line-level annotations. It contains 1000 training images for training and 500 images for testing.

4.3. Ablation Study

In this section, we utilize the point detector for text center point prediction to conduct the ablation studies. It is worth noting that the recognizer in this paper is trained in the word-level synthetic recognition data. Therefore, the performance on CTW1500 will be lower than the word-level benchmark ICDAR2015 and TotalText.



Figure 4. Visualization results of training using the generated polygons. The green-filled polygons are the ground truths, the blue-filled polygons are the ground truths that are marked as “don’t care”, and the red polygons are the detection results. Method (GT) represents training using the polygons generated from GT points; while Method (Ours) represents using the polygons generated from existing single-point text spotter (SPTS v2). Best view in screen.

Method Selection			ICDAR15			TotalText			CTW1500		
AGM	PGM	PRM	Precision	Recall	Hmean	Precision	Recall	Hmean	Precision	Recall	Hmean
✓	✓	✓	77.1	70.7	73.8	72.0	71.0	71.5	54.2	56.7	55.4
	✓	✓	18.5	17.1	17.8	11.8	11.5	11.6	6.6	10.3	8.1
✓		✓	70.7	65.2	67.9	60.7	60.2	60.4	49.1	51.4	50.2
✓	✓		74.4	68.3	71.2	71.7	70.6	71.1	54.0	56.4	55.2

Table 1. Ablation studies on the proposed designs. AGM means the Anchor Generation Module. PGM means the Polygon Generation Module. PRM means the Polygon Rectification Module.

Ablation study of AGM. To evaluate the efficacy of the Anchor Generation Module (AGM), we conducted a comparative analysis between randomly selected anchors and those generated using AGM. The results are presented in Tab. 1. A comparison between the results in the first and second rows underscores the pivotal role played by AGM in polygon generation. The absence of AGM results in a performance decline of 56.0% on IC15, 59.9% on TotalText, and 47.3% on CTW1500, respectively. AGM demonstrates its capability by generating a coarse anchor, providing the width and height of the text for the subsequent modules based on recognition confidence. The absence of these coarse anchors supplied by AGM renders the PGM incapable of generating precise boundary points. The qualitative results are illustrated in Fig. 5.

Ablation study of PGM. To verify the effectiveness of the Polygon Generation Module (PGM), we conducted an

ablation study where we removed the PGM and directly forwarded the coarse anchors to the Polygon Rectification Module. As illustrated in Tab. 1, without PGM, the performance decreased by approximately 5.9% on the IC15 dataset, 11.1% on the TotalText dataset, and 5.2% on the CTW1500 dataset, respectively. The qualitative results are illustrated in Fig. 5. The PGM serves to transform the coarse anchors into boundary points that are in closer proximity to the text region, which plays a crucial role of the PGM in improving the quality of the boundary points. Notably, the PGM is supervised by the recognition loss without the detection annotations.

Ablation study of PRM. We conducted a comparative analysis to evaluate the impact of the Polygon Rectification Module (PRM) in Tab. 1. The PRM notably influences IC15, revealing that boundary points from PGM often extend beyond the boundaries in multi-oriented text.

Attn. Threshold	ICDAR2015			TotalText			CTW1500		
	Precision	Recall	Hmean	Precision	Recall	Hmean	Precision	Recall	Hmean
0.05	77.1	70.7	73.8	71.7	70.7	71.2	53.5	56.0	54.7
0.04	77.1	70.6	73.7	71.7	70.7	71.2	53.6	56.1	54.8
0.03	76.9	70.5	73.6	71.9	70.9	71.4	53.7	56.1	54.9
0.02	76.9	70.5	73.6	71.9	70.9	71.4	53.8	56.3	55.0
0.01	76.0	69.7	72.7	72.0	71.0	71.5	54.2	56.7	55.4

Table 2. The influence of attention threshold. Attn. represents attention.

Threshold	Results Generated from GT Points								
	ICDAR2015			TotalText			CTW1500		
IOU	Precision	Recall	Hmean	Precision	Recall	Hmean	Precision	Recall	Hmean
0.1	97.1	89.8	93.3	97.6	92.8	95.1	94.8	90.6	92.7
0.3	95.7	88.6	92.0	91.4	86.6	88.8	80.0	76.4	78.1
0.5	85.3	79.0	82.0	76.0	72.3	74.1	62.8	60.0	61.3
0.7	48.4	44.8	46.6	32.9	31.3	32.1	27.9	26.7	27.3
Results Generated from SPTS v2 Points									
IOU	Precision	Recall	Hmean	Precision	Recall	Hmean	Precision	Recall	Hmean
0.1	92.5	84.8	88.5	92.0	90.7	91.4	85.8	89.7	87.7
0.3	90.0	82.5	86.1	87.9	86.7	87.3	71.8	75.0	73.4
0.5	77.1	70.7	73.8	72.0	71.0	71.5	54.2	56.7	55.4
0.7	32.8	30.1	31.4	27.3	26.9	27.1	23.3	24.4	23.8
Results of Original SPTS v2									
DIST	Precision	Recall	Hmean	Precision	Recall	Hmean	Precision	Recall	Hmean
5	59.0	54.4	57.0	46.1	45.4	45.7	61.8	63.7	62.7
10	88.4	80.3	84.2	70.0	69.1	69.5	74.3	76.5	75.7
20	92.7	84.2	88.2	82.7	81.6	82.1	83.0	85.4	84.2
30	94.0	85.4	89.5	87.8	86.6	87.2	85.3	87.9	86.6

Table 3. The comparison involves SPTS v2 point detection and Point2Polygon predictions. 'GT Points' indicates polygons created from Ground Truth points, and 'SPTS v2 Points' denotes polygons derived from SPTS v2 predicted points.

The qualitative results are illustrated in Fig. 5. Without the PRM, the boundary points generated by the PGM tend to excessively broaden horizontally, extending well beyond the actual text region. However, the implementation of the PRM effectively eliminates these surplus points beyond the text’s boundaries.

Ablation study of the attention threshold in PRM. Our method presets the attention threshold to remove the excess points beyond the text’s range. To evaluate the influence of different thresholds, we conduct experiments while maintaining the model structure unchanged, only modifying the threshold. The results are presented in Tab. 2. Notably, as the threshold increased, we observed a gradual improvement in the performance of IC15. Conversely, for Totaltext and CTW1500, we observed a slight decline in performance

with higher attention thresholds. This suggests that there is a greater need for rectification in the case of multi-oriented text, which should follow high-quality standards.

4.4. Quality of the Generated Polygon

In this section, we conduct experiments to verify the quality of polygons generated by our methods. There are two types of points considered: 1) Center points generated based on the Ground Truth (GT Points) of the testing set; and 2) Center points predicted using SPTS v2 (SPTS v2 Points). Subsequently, we generate polygons based on these two types and evaluate their performance. It is important to note that the accuracy of polygons generated from SPTS v2 is constrained by its performance.

We test the performance across four IOU thresholds. As illustrated in Tab. 3, under 0.5 IOU threshold, polygons generated from GT Points achieve 82% on IC15, 74.1% on To-

Model	ICDAR2015			TotalText			CTW1500		
	Precision	Recall	Hmean	Precision	Recall	Hmean	Precision	Recall	Hmean
Detection Results Trained Using Original Ground Truth									
DBNet + resnet50	89.8	79.4	84.3	84.5	79.6	82.0	68.3	63.4	65.8
PSENet + resnet50	83.6	76.2	79.7	82.3	77.2	79.6	77.5	82.9	80.1
PANet + resnet18	83.1	72.9	77.6	85.9	77.3	81.4	86.3	82.0	84.1
Detection Results Trained Using Polygons Generated from SPTS v2 Point									
DBNet + resnet50	82.4	67.9	77.5	81.8	65.0	72.4	61.5	42.8	50.5
PSENet + resnet50	77.4	54.2	63.8	76.5	55.8	64.5	56.3	41.7	47.9
PANet + resnet18	67.8	51.1	58.3	77.9	55.9	65.1	67.2	42.9	52.3
Detection Results Trained Using Polygons Generated from GT Point									
DBNet + resnet50	87.4	65.5	74.9	84.6	65.9	74.1	62.9	44.1	51.8
PSENet + resnet50	72.7	58.2	64.6	76.7	56.9	65.3	56.7	38.6	46.0
PANet + resnet18	77.2	51.8	62.0	79.1	57.4	66.5	58.0	45.5	51.0

Table 4. Detection results using the polygon generated from our method. All experiments are conducted in the mmocr. The Raw Ground Truth means using the polygon from Ground Truth to train the detector. SPTS v2 Point Generated Annotation involves utilizing points from SPTS v2 to generate polygons for training the detector. GT Point Generated Annotation entails using points from the ground truth (GT) to generate polygons for training the detector.



Figure 5. Qualitative results of AGM (left), PGM (middle) and PRM (right). The upper and lower line shows results with or without using the module, respectively. Best view in screen.

talText, and 61.3% on CTW1500, respectively, in terms of Hmean; while the polygons generated from SPTS v2 Points achieve 73.8% on IC15, 71.5% on TotalText, and 55.4% on CTW1500, respectively. We can also observe that as the IOU threshold decreases, the performance gap between SPTS v2 Points and GT Points narrows.

To further evaluate the performance of the generated points from SPTS v2, we also directly compare the detection performance with SPTS v2. The term “DIST” denotes the number of pixels within which predicted points are considered accurate relative to ground truth points. To relate IOU to DIST, we initially calculate the difference in performance between polygons generated from SPTS v2 Points

and GT Points on IC15 under a 0.5 IOU. Using SPTS v2 Points achieves 89.5% of the performance achieved using GT Points. Therefore, at an IOU of 0.5, the corresponding DIST is about 30. Mapping this to IC15, the generated polygons from SPTS v2 Points achieve 82.5% of SPTS v2’s performance. The results demonstrate our method can be effectively integrated into an existing single-point text spotter to produce polygon results without increasing the labeling costs.

4.5. Evaluating on Existing Detector

To facilitate a comprehensive comparison between the polygons generated by our method and the ground truth polygons, we utilized these polygons for training the DBNet [17], PANet [33], and PSENet [32] models. Specifically, we employed quadrilaterals as ground truth annotations for the IC15 dataset, polygons for the TotalText dataset and the CTW1500 dataset. All training procedures were conducted within the mmocr framework [15], and the results are outlined in Tab. 4 with an IoU threshold of 0.5. The results are categorized into two types: 1) GT Points Generated Annotation, where we generated the center points based on the ground truth (GT Points) from the training set, and used these polygons to train the detection model; 2) SPTS v2 Points Generated Annotation, where we utilized the SPTS v2 to predict the center points (SPTS v2 Points), and employed the resulting polygons for training the detection model.

By employing the SPTS v2 Points Generated Annotation for training the detector, the DBNet demonstrates notable performance with an accuracy of 77.5% on IC15, 72.4% on TotalText, and 50.5% on CTW1500. Similarly, the PSENet achieves 63.8% on IC15, 64.5% on TotalText, and 47.9% on

CTW1500, while the PANet attains 58.3% on IC15, 65.1% on TotalText, and 52.3% on CTW1500. Furthermore, it is evident from the results that SPTS v2 Points Generated Annotation achieves approximately 86% of the performance observed when using polygons from the ground truth to train the model. This highlights the effectiveness of our method. Further comparison between GT Points Generated Annotation and SPTS v2 Points Generated Annotation reveals that our method achieves superior performance when training with SPTS points in certain cases compared to GT points. This implies that our method can serve as an automatic labeling tool when coupled with a point detector. Qualitative results are shown in the Fig. 4.

Moreover, by comparing the seventh row of Tab. 3 with the fourth row of Tab. 4, we observe that the results of DBNet using polygons generated from SPTS v2 Point outperform the performance achieved with polygons without training by detectors, with 3.7% and 0.9% improvement on IC15 and Totaltext, respectively, in terms of Hmean. It demonstrates that while the generated polygon may not achieve perfect accuracy, its overall quality remains high. This high-quality polygon allows the detector to further exceed the performance of the polygon generated from SPTS v2 Point.

5. Discussion

5.1. Visualization

Example visualization results of the polygon generated by Point2Polygon are presented in Fig. 6. Point2Polygon can generate visually-plausible polygons for various of text including horizontal, multi-oriented, and arbitrarily-shaped scene text. We also observe it shows robust performance in generating polygons for slightly curved text.

5.2. Limitation

A noted limitation is that the generated polygon does not entirely align with the Ground Truth, especially for some highly curved texts, as shown in the last line of Fig. 6. Nevertheless, we believe that the current level of precision has the potential to alleviate the annotation labor and serve as a foundation to drive future studies in this domain. Additionally, studies such as TTS [14], SPTS [24], and TOSS [29] have shown that without any specific location information but with transcription-only, successfully training a robust scene text spotter is still viable. This suggests that spotters are able to implicitly learn text locations in images. However, Point2Polygon currently requires a reference point to generate the polygon, indicating room for potential enhancement.



Figure 6. Visualization results of Point2Polygon. The first row shows horizontal examples. The second row shows tilted examples. The third row shows slightly curved examples. The fourth row shows some failure examples. Best view in screen.

6. Conclusion

We introduce the Point2Polygon for effectively evolve points into polygons from coarse to fine through multi-granularity recognition information. A key aspect of our method is the use of synthetic datasets for recognition information, effectively bypassing the need for any manual annotation beyond a single point. Our extensive experiments demonstrate the strength and reliability of our method. These tests involve direct comparisons to ground truth data, applying our approach in conjunction with existing text detection systems, and assessing its performance when integrated with the single-point spotting method. While there is a noticeable distinction between our generated results and actual ground truths, Point2Polygon has shown encouraging results. Additionally, it is worth noting that this method may aid in automatically producing more detailed information in large-scale image-text pairs, thereby assisting in extensive training. We believe this approach not only demonstrates significant potential in the realm of point-to-polygon evolution but also sets a strong baseline for future research in this area.

References

- [1] Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoon Yun, and Hwalsuk Lee. Character region awareness for text detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9365–9374, 2019. 1
- [2] Fred L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(6):567–585, 1989. 3
- [3] Zhanzhan Cheng, Yangliu Xu, Fan Bai, Yi Niu, Shiliang Pu, and Shuigeng Zhou. Aon: Towards arbitrarily-oriented text recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5571–5579, 2018. 2
- [4] Chee-Kheng Ch’ng, Chee Seng Chan, and Cheng-Lin Liu. Total-text: toward orientation robustness in scene text detection. *International Journal on Document Analysis and Recognition (IJ DAR)*, 23(1):31–52, 2020. 4
- [5] Yongkun Du, Zhineng Chen, Caiyan Jia, Xiaoting Yin, Tianlun Zheng, Chenxia Li, Yuning Du, and Yu-Gang Jiang. Svtr: Scene text recognition with a single visual model. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 884–890. International Joint Conferences on Artificial Intelligence Organization, 2022. Main Track. 4
- [6] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2315–2324, 2016. 4
- [7] Wenhao He, Xu-Yao Zhang, Fei Yin, and Cheng-Lin Liu. Deep direct regression for multi-oriented scene text detection. In *Proceedings of the IEEE international conference on computer vision*, pages 745–753, 2017. 2
- [8] Han Hu, Chengquan Zhang, Yuxuan Luo, Yuzhuo Wang, Junyu Han, and Errui Ding. Wordsup: Exploiting word annotations for character based text detection. In *Proceedings of the IEEE international conference on computer vision*, pages 4940–4949, 2017. 1, 2
- [9] Mingxin Huang, Yuliang Liu, Zhenghao Peng, Chongyu Liu, Dahua Lin, Shenggao Zhu, Nicholas Yuan, Kai Ding, and Lianwen Jin. Swintextspotter: Scene text spotting via better synergy between text detection and text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4593–4603, 2022. 2
- [10] Mingxin Huang, Jiaxin Zhang, Dezhi Peng, Hao Lu, Can Huang, Yuliang Liu, Xiang Bai, and Lianwen Jin. ES-TextSpotter: Towards better scene text spotting with explicit synergy in transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19495–19505, 2023. 2
- [11] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*, 2014. 4
- [12] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 1156–1160. IEEE, 2015. 4
- [13] Taeho Kil, Seonghyeon Kim, Sukmin Seo, Yoonsik Kim, and Daehee Kim. Towards unified scene text spotting based on sequence generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15223–15232, 2023. 1, 3
- [14] Yair Kittenplon, Inbal Lavi, Sharon Fogel, Yariv Bar, R Manmatha, and Pietro Perona. Towards weakly-supervised text spotting using a multi-task transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4604–4613, 2022. 1, 2, 8
- [15] Zhanghui Kuang, Hongbin Sun, Zhizhong Li, Xiaoyu Yue, Tsui Hin Lin, Jianyong Chen, Huaqiang Wei, Yiqin Zhu, Tong Gao, Wenwei Zhang, et al. Mmocr: A comprehensive toolbox for text detection, recognition and understanding. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3791–3794, 2021. 7
- [16] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *Thirty-first AAAI Conference on Artificial Intelligence*, 2017. 2
- [17] Minghui Liao, Zhaoyi Wan, Cong Yao, Kai Chen, and Xiang Bai. Real-time scene text detection with differentiable binarization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 11474–11481, 2020. 7
- [18] M. Liao, P. Lyu, M. He, C. Yao, W. Wu, and X. Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2):532–548, 2021. 2
- [19] Chongyu Liu, Yuliang Liu, Lianwen Jin, Shuaitao Zhang, Canjie Luo, and Yongpan Wang. EraserNet: End-to-end text removal in the wild. *IEEE Transactions on Image Processing*, 29:8760–8775, 2020. 1
- [20] Ruijin Liu, Ning Lu, Dapeng Chen, Cheng Li, Zejian Yuan, and Wei Peng. Pbformer: Capturing complex scene text shape with polynomial band transformer. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 2112–2120, 2023. 2
- [21] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European Conference on Computer Vision*, pages 21–37. Springer, 2016. 2
- [22] Yuliang Liu, Lianwen Jin, Shuaitao Zhang, Canjie Luo, and Sheng Zhang. Curved scene text detection via transverse and longitudinal sequence connection. *Pattern Recognition*, 90: 337–345, 2019. 4
- [23] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9809–9818, 2020. 2
- [24] Yuliang Liu, Jiaxin Zhang, Dezhi Peng, Mingxin Huang, Xinyu Wang, Jingqun Tang, Can Huang, Dahua Lin, Chun-

- hua Shen, Xiang Bai, et al. Spts v2: Single-point scene text spotting. *arXiv preprint arXiv:2301.01635*, 2023. 1, 2, 8
- [25] Canjie Luo, Lianwen Jin, and Zenghui Sun. Moran: A multi-object rectified attention network for scene text recognition. *Pattern Recognition*, 90:109–118, 2019. 4
- [26] Dezhi Peng, Xinyu Wang, Yuliang Liu, Jiabin Zhang, Mingxin Huang, Songxuan Lai, Jing Li, Shenggao Zhu, Dahua Lin, Chunhua Shen, et al. Spts: single-point text spotting. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4272–4281, 2022. 2
- [27] Yadong Qu, Qingfeng Tan, Hongtao Xie, Jianjun Xu, Yuxin Wang, and Yongdong Zhang. Exploring stroke-level modifications for scene text editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2119–2127, 2023. 1
- [28] Baoguang Shi, Mingkun Yang, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Aster: An attentional scene text recognizer with flexible rectification. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2035–2048, 2018. 2, 4
- [29] Jingqun Tang, Su Qiao, Benlei Cui, Yuhang Ma, Sheng Zhang, and Dimitrios Kanoulas. You can even annotate text with voice: Transcription-only-supervised text spotting. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4154–4163, 2022. 2, 8
- [30] Shangxuan Tian, Shijian Lu, and Chongshou Li. Wetext: Scene text detection under weak supervision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1492–1500, 2017. 1, 2
- [31] Zhi Tian, Weilin Huang, Tong He, Pan He, and Yu Qiao. Detecting text in natural image with connectionist text proposal network. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part VIII 14*, pages 56–72. Springer, 2016. 2
- [32] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9336–9345, 2019. 2, 7
- [33] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8440–8449, 2019. 2, 7
- [34] Yuxin Wang, Hongtao Xie, Shancheng Fang, Yadong Qu, and Yongdong Zhang. Pert: a progressively region-based network for scene text removal. *arXiv preprint arXiv:2106.13029*, 2021. 1
- [35] Liang Wu, Chengquan Zhang, Jiaming Liu, Junyu Han, Jingtuo Liu, Errui Ding, and Xiang Bai. Editing text in the wild. In *Proceedings of the 27th ACM international conference on multimedia*, pages 1500–1508, 2019. 1
- [36] Weijia Wu, Enze Xie, Ruimao Zhang, Wenhai Wang, Ping Luo, and Zhou Hong. Polygon-Free: Unconstrained scene text detection with box annotations. In *Proceedings of the IEEE International Conference on Image Processing*, pages 1226–1230, 2022. 1, 2
- [37] Maoyuan Ye, Jing Zhang, Shanshan Zhao, Juhua Liu, Tongliang Liu, Bo Du, and Dacheng Tao. DeepSolo: Let transformer decoder with explicit points solo for text spotting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19348–19357, 2023. 2
- [38] Matthew D Zeiler. Adadelta: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012. 4
- [39] Mengbiao Zhao, Wei Feng, Fei Yin, and Cheng-Lin Liu. Texts as points: Scene text detection with point supervision. *Pattern Recognition Letters*, 170:1–8, 2023. 1
- [40] Tianlun Zheng, Zhineng Chen, Shancheng Fang, Hongtao Xie, and Yu-Gang Jiang. Cdistnet: Perceiving multi-domain character distance for robust text recognition. *International Journal of Computer Vision*, pages 1–19, 2023. 4
- [41] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 5551–5560, 2017. 2