Towards robust curve text detection with conditional spatial expansion

Liu, Zichuan; Lin, Guosheng; Yang, Sheng; Liu, Fayao; Lin, Weisi; Goh, Wang Ling

2019

Liu, Z., Lin, G., Yang, S., Liu, F., Lin, W., & Goh, W. L. (2019). Towards robust curve text detection with conditional spatial expansion. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). doi:10.1109/CVPR.2019.00744

<https://hdl.handle.net/10356/144281>

<https://doi.org/10.1109/CVPR.2019.00744>

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at:
<https://doi.org/10.1109/CVPR.2019.00744>

Downloaded on 26 Apr 2023 18:06:42 SGT

Towards Robust Curve Text Detection with Conditional Spatial Expansion

Zichuan Liu¹, Guosheng Lin¹, Sheng Yang¹, Fayao Liu², Weisi Lin¹ and Wang Ling Goh¹

¹Nanyang Technological University, Singapore

²University of Adelaide, Australia

{zliu016, syang014}@e.ntu.edu.sg, {gslin, wslin, ewlgoh}@ntu.edu.sg, fayao.liu@gmail.com

Abstract

It is challenging to detect curve texts due to their irregular shapes and varying sizes. In this paper, we first investigate the deficiency of the existing curve detection methods and then propose a novel Conditional Spatial Expansion (CSE) mechanism to improve the performance of curve text detection. Instead of regarding the curve text detection as a polygon regression or a segmentation problem, we treat it as a region expansion process. Our CSE starts with a seed arbitrarily initialized within a text region and progressively merges neighborhood regions based on the extracted local features by a CNN and contextual information of merged regions. The CSE is highly parameterized and can be seamlessly integrated into existing object detection frameworks. Enhanced by the data-dependent CSE mechanism, our curve text detection system provides robust instance-level text region extraction with minimal post-processing. The analysis experiment shows that our CSE can handle texts with various shapes, sizes, and orientations, and can effectively suppress the false-positives coming from text-like textures or unexpected texts included in the same RoI. Compared with the existing curve text detection algorithms, our method is more robust and enjoys a simpler processing flow. It also creates a new state-of-art performance on curve text benchmarks with F-score of up to 78.4%.

1. Introduction

In recent years, great progress has been made in text detection. The performance has been enhanced by the advanced object detection and segmentation frameworks based on Neural Networks. Although detecting words or text lines with different sizes and orientations has been well tackled by recently proposed methods [28, 24, 34, 32, 23], detecting curve texts remains a challenging problem.

The main challenges of curve text detection come from irregular shapes and highly varying orientations. The traditional bounding box representation does not scale well

Figure 1. Problems of existing curve text detection methods: Two region proposals shown in (a) and (d) inevitably include unexpected texts since they are closed to each other. Thus, it causes failures for Mask RCNN based methods and polygon regression methods demonstrated in (b) and (e). Our CSE method demonstrated in (c) and (f) shows strong robustness to this situation and brings significant performance gain.

in the curve scenario since one box may cover multiple text objects. Therefore, the recently proposed curve text detection algorithms [39, 27, 40] follow a two-stage detect-and-refine approach to generate elaborated polygons or boundaries. In these methods, a CNN based text detector is applied to locate the regions containing texts, and then a segmentation or polygon regression algorithm is performed on these regions to produce a tight polygon or boundary. Both methods highly depend on the accurate region proposal provided by the text detector. They prefer a proposed region with only one targeted object included which reduces the ambiguity of the sampled features. Although an oriented box regression is preferable, it often fails in the curve text scenario [29, 14]. The recently proposed curve text detection methods turn to predict horizontal rectangles in the first stage, which is inevitably affected non-targeted texts in the sampled regions.

Specifically, both segmentation and regression based

methods can be disrupted by the unexpected texts included in the same box region. The segmentation based methods can fail to distinguish the targeted text from the others and misclassify the unexpected texts or text-like patterns as positive, shown in Fig. 1 (b). The regression based methods can produce incorrect boundaries by indistinguishably considering all texts as one object. Moreover, as shown in Fig. 1 (e), the regression results (produced by predicting an offset of a proposed box region) are highly coupled with previous stage box proposals. A poor box proposal greatly affects the final polygon which degrades the performance.

To tackle the problem mentioned above, we propose a novel Conditional Spatial Expansion (CSE) mechanism, which acts as a second-stage component applied in the widely adopted two-stage detection workflow. Our method is derived from conditional modeling of dependency between an interior point (called a seed) and the rest parts of a text instance. The CSE can be seen as a conditional prediction process which retrieves an instance level text region by seeding and expanding. Starting with an arbitrary interior point (seed) of a text region, CSE selectively expands its territory based on location observation of the image patches and the context inferred from merged regions. Compared with the segmentation based method, our CSE is extremely discriminative especially when texts are close to each other as demonstrated in Fig. 1 (c). It provides a controllable approach to extract an expected text region with minimum efforts of post-processing. On the other hand, our CSE is highly flexible since a seed can be specified at any location within the targeted text region. Compared with the polygon regression methods, the seeding-and-expanding paradigm has less coupling with the previous text detector. Base on a coarse region proposal, our CSE is directly applied on the spatial features produced by the backbone, which preserves all the spatial information and will not be affected by the imperfect region proposals, shown in Fig. 1 (f). The experiments show that our method outperforms the existing curve text detection methods on public benchmarks. The contributions of this work are summarized as follows:

- The curve text detection is formulated as a conditional region expansion problem, which initializes a seed within a text region and then progressively retrieves targeted object by region expansion;
- The spatial dependency between the seed and the rest part of an object is modeled by a parameterized Conditional Spatial Expansion mechanism, which allows us to selectively extract a text region indicated by a seed with high area precision;
- Our CSE acts as a second-stage text extractor which can be seamlessly integrated into existing object detection workflows;
- The arbitrariness of seed's location and high spatial selectivity of our method reduce coupling with the previous detector and thus provide flexible and robust boundary prediction;
- Our method outperforms the existing curve text detection methods on public curve text datasets with F-measurement of 80.2% on Total-Text [5] and 78.4% on CTW-1500 [39].

2. Related Works

2.1. Quadrilateral Text Detection

In the quadrilateral text detection, the ground-truths are constrained to a rectangle or a quadrilateral. Base on the types of targets to be retrieved, text detection methods can be categorized as detection based methods and segmentation based methods.

The detection based method follows the object detection frameworks [9, 31, 20, 30, 21, 10] driven by Convolutional Neural Networks (CNNs) [17]. TextBoxes [18] adopts SSD as a base detector and handles variation of aspect ratios of text instance by elaborated reference box design. As the variants of the faster RCNN [30], the Rotation Region Proposal Network (RRPN) [29] and Rotational Region CNN (R²CNN) [14] are designed to detect arbitrarily oriented texts in a two-stage manner. In addition, EAST [43] and DeepReg [11] are proposed to directly regress the geometry of a text instance.

The segmentation based methods are mostly designed to extract long text lines in an image. They interpret text detection as a semantic segmentation problem which has been well addressed by the Fully Convolutional Neural Networks (FCNs) [25, 7, 8, 41]. Zhang *et al.* [42] combines FCN and MSER [13] to recognize text blocks and then extract corresponding characters. Yao *et al.* [38] applies FCN to predict multiple properties of texts, such as text regions and orientations, to extract the target text regions. To distinguish adjacent text instances, the component segmentation method [24, 23, 35, 6, 36] is proposed, where a text region is broken into several components, which will be combined into different instances by data-driven clustering, inter-node communication or post-processing.

2.2. Curve Text Detection

Although the methods reviewed above have succeeded in the quadrilateral text detection, most of them cannot scale well in the case of arbitrary text shape. New representations and detection framework are proposed for this task. Liu *et al.* [39] propose Transverse and Longitudinal Offset Connection (TLOC) method based on Faster RCNN and Recurrent Neural Networks (RNNs) to directly regress the polygon shape of text regions.

Mask Text-Spotter [27] regards the curve text detection as an instance segmentation problem and applies the Mask-RCNN to produce boundaries of text instances. TextSnake [26] adopts FCN as the base detector and extract text instance by detecting and assembling local components.

Most of the existing curve text detection methods are potentially modeling the posterior probability between observed image patches and the ground-true foreground labels. Instead, our method captures the dependency between arbitrary image patches and the rest of the text regions. The proposed modeling is naturally robust to ambiguity caused by multiple text instances included in one RoI. Moreover, our CSE considers more local details and thus can produce more elaborated text boundaries.

3. Method

3.1. Overview

Our method retrieves an instance level text region by seeding and then expanding. A seed uniquely indicates an object and can be arbitrarily initialized at the interior of an object region. Seeding is to select a location within an object from which the corresponding object region is extracted by expanding. As will be illustrated in Sect. 3.4, a seed is initialized by an object detector with a predicted box center. Starting with a seed, the expanding is conducted by selectively merging adjacent sub-regions to form a targeted object region. As shown in Fig. 2, the sub-regions are abstracted as feature points or nodes, which are sampled from the input image at discrete locations. They are organized as a grid and locally assigned an expanding indicator $y \in \mathbb{R}^5$ to represent the merging direction to neighborhood nodes. Five entries of y denote the probabilities of all possible merging directions (*to – bottom*, *to – right*, *to – left*, *to – top* and *none*). A node will be merged into an existing object region indicated by a seed if its major merging direction is pointing to its neighborhoods that already belong to the object region. An instance level object boundary can be easily produced by mapping all the positive nodes back to the original image and extracting the contour.

3.2. Modeling

The seeding-and-expanding paradigm provides a flexible and controllable way to extract object regions with minimal post-processing. It also reduces the performance coupling with the first-stage detector by allowing arbitrary initialization of a seed node. However, with different seed locations, the dynamics of the region expansion are different. For a specific node, the state of expanding indicator varies with the seed’s location and also depends on the expanding indicators of its neighborhoods. Therefore, obtaining the expanding indicators should be regarded as a

Figure 2. Our CSE inside-out explores the expanding indicator of every node and merges nodes with merging direction pointing to nodes that are already in the object region.

conditional prediction problem. Centered at a seed shown in Fig. 2, we divide the region into sections (defined as nodes with the same minimum number of steps to a seed) using a set of contours. We assume that the expanding indicators of nodes P_k within the k -th section are mutually independent and their states depend only on the current spatial feature X_k and the states of nodes within previous sections $P_{k-1}, P_{k-2}, \dots, P_0$. The optimal estimator of Y_k can be represented by

$$\hat{Y}_k = \underset{Y_k}{\operatorname{argmax}} \Pr(Y_k | Y_{k-1}, \dots, Y_0; X_k, \dots, X_0), \quad (1)$$

which maximizes the posterior probability of Y_k when observing spatial features $X_{(\cdot)} := \{x(p) | p \in P_{(\cdot)}\}$ and indicators of previous nodes $Y_{(\cdot)} := \{y(p) | p \in P_{(\cdot)}\}$. This conditional modeling allows the region expansion to be adaptable with an arbitrarily initialized seed location. Also, it effectively differentiates expected object from the others by considering the context derived from a seed. Moreover, independence assumption among node in the same section results in a dendritic Conditional Spatial Expansion process with high-level parallelism.

3.3. Conditional Spatial Expansion

To estimate conditional probability illustrated in Eq. 1, we develop a highly parameterized Condition Spatial Expansion (CSE) mechanism. Given a seed inside an object region, we construct its neighborhood feature points by expanding a $S \times S$ grid and then sampling the features produced by the backbone network using bilinear interpolation, as shown in Fig. 3 (a). Starting with the seed node, our CSE explores every single node inside-out and computes corresponding y and the transition vectors $h_o = [h_o^b, h_o^r, h_o^l, h_o^t]^T$ to its neighborhoods based on the current sampled feature $x \in \mathbb{R}^{d_x}$, local state $c \in \mathbb{R}^d$

Figure 3. Given a seed located at the interior of a text region, we expand a grid with $S \times S$ points and sample the feature produced by backbone at these locations using bilinear interpolation. The CSE computation starts with the seed and spreads to the adjacent feature nodes. Each node takes the outputs and the hidden states of previous nodes as input and produces new hidden state and output.

and transition vectors $h_i^b, h_i^r, h_i^l, h_i^t \in \mathbb{R}^d$ coming from the adjacent feature points. The transition vectors encode the position sensitive information which helps the CSE to be aware of the relative location of the current node to the seed. Depending on the relative position to the seed, the inputs and outputs for the nodes are illustrated in Fig. 3 (c). For a node in P_k , our CSE only takes the c and h as inputs from P_{k-1} and output new h to P_{k+1} . This constructs an inference process originated from the seed which propagates the contextual information among the grid in a dendritic manner. The computation of the nodes in the same section is independent and thus can be fully parallelized on GPUs. The computation complexity is linear to the side of the grid, which is computationally efficient.

Inside a specific node, the computation is illustrated by a computing graph shown in Fig 4. Without loss of generality, we denote all possible inputs from neighborhood nodes by $c_i \in \mathbb{R}^{4d}$, $h_i \in \mathbb{R}^{5d}$, and $y_i \in \mathbb{R}^{20}$, which are represented by

$$c_i = [c_i^b, c_i^r, c_i^l, c_i^t]^T, \quad (2)$$

$$h_i = [h_i^c, h_i^b, h_i^r, h_i^l, h_i^t]^T, \quad (3)$$

$$y_i = [y_i^b, y_i^r, y_i^l, y_i^t]^T, \quad (4)$$

where $c_i^{(\cdot)} \in \mathbb{R}^d$, $h_i^{(\cdot)} \in \mathbb{R}^d$ and $y_i^{(\cdot)} \in \mathbb{R}^5$ denotes the local states, transition vectors and the expanding indicators of neighborhood nodes¹. Here, only parts of c_i , h_i and y_i

are available according to the relative position to the seed node², and the rest will be set to zeros. Particularly, h_i^c is defined to uniquely indicate the seed node. Other than h_i^c of a seed which is learned by backpropagation, we explicitly set the h_i^c of other node to zeros.

From the current observed feature x , transition vectors h_i and the predicted expanding indicators of neighborhoods y_i , we compute a candidate local state $\tilde{c}$ by

$$\tilde{c} = \tanh(W_c \times s + b_c), \quad (5)$$

$$s = [x, y_i, h_i]^T, \quad (6)$$

where $\times$ represents the matrix multiplication, $W_c \in \mathbb{R}^{d \times (d_x + 5d + 20)}$ and $b_c \in \mathbb{R}^d$ denote weights and bias of linear transform before a tanh activation. We apply the gating mechanism [12] to combine the local state c_b , c_r , c_l and c_t from the neighborhoods with the current candidate state $\tilde{c}$ to obtain the local state of current node c , which is formulated as

$$c = \delta(c_b \cdot g_c^b + c_r \cdot g_c^r + c_l \cdot g_c^l + c_t \cdot g_c^t + \tilde{c} \cdot g_{\tilde{c}}), \quad (7)$$

where δ denotes the layer normalization operator [4], $\cdot$ is the element-wise multiplication, and g_b , g_r , g_l , g_t , $g_{\tilde{c}}$ represent the outputs of gating function which can be further illustrated by

$$g_c = [g_{c_i}^b, g_{c_i}^r, g_{c_i}^l, g_{c_i}^t]^T = \sigma(W_{g_{c_i}} \times s + b_{g_{c_i}}), \quad (8)$$

$$g_{\tilde{c}} = \sigma(W_{g_{\tilde{c}}} \times s + b_{g_{\tilde{c}}}). \quad (9)$$

¹The superscripts indicate the relative position to the current node, b-bottom, r-right, l-left, t-top and c-current.

²The non-zero inputs are indicated in Fig. 3 (c)

Figure 4. Computing Graph inside a Node.

Here, $W_{g_{c_i}}$, $b_{g_{c_i}}$ and $W_{g_{\tilde{c}}}$, $b_{g_{\tilde{c}}}$ are defined as the weight matrix and bias to map s into corresponding gating vectors g_{c_i} and $g_{\tilde{c}}$. Since the local state c is essentially the weighted sum of state values of the previous nodes, the values of c increases exponentially with k in our two-dimensional scenario. This significantly harms the numerical stability in both training and testing phases. Thus, the layer normalization technique is essential for the CSE to ensure the convergence of training and prevent overflow in testing. Finally, the expanding indicator y and the output transition vectors h_o are derived from the local state c , which are illustrated as follows

$$g_o = \sigma(W_{g_o} \times s + b_{g_o}), \quad (10)$$

$$h_o = [h_o^b, h_o^r, h_o^l, h_o^t]^T = \tanh(c) \cdot g_o + b_o, \quad (11)$$

$$y = \text{softmax}(W_y \times c + b_y), \quad (12)$$

where W_{g_o} and b_{g_o} represent the weight matrix and bias used to produce the corresponding gating signal, and W_y and b_y transform c into logits before feeding to *softmax* activation.

3.4. Seeding

In our CSE method, a seed is assumed to be located within an object region. This prerequisite can be easily guaranteed by using the outputs of an off-the-shelf object detector. Here, we adopt detected box centers and shapes by Faster RCNN [31] to decide seed locations and the shapes, and uniformly sample $S \times S$ features using bilinear interpolation from a region indicated by a bounding box. In fact, our CSE method only requires a seed to be located within an object region with sampling grid coarsely covering the targeted object. Moreover, given a sampling grid, any node within the object region can be specified as a seed. As shown in Sect. 4.2, randomly initializing seed location and corresponding grid size does not significantly affect the performance. Therefore, a weaker detector, which

Figure 5. Ground-truth Labeling: All arrows represent the candidate merging directions for nodes, and the corresponding merging directions predicted by CSE are marked. The candidate directions with highest scores are label as positive directions, which are represented by solid arrows.

is easy to optimize, could be sufficient for CSE to produce satisfactory results.

3.5. Optimization

Labeling In the training phase, the ground-true merging directions are labeled using the strategy illustrated in Fig. 5. For each grid in CSE, we first label the nodes within the target ground-true object region as positive and the rest as negative. For every positive node, we search its neighborhood positive nodes at the previous section and label the corresponding merging directions as candidate directions³. Among the candidate directions of the same node, we only label the one with the highest score as the final positive merging direction. For the seed node, we always label its y_s as positive.

Loss Function We apply cross-entropy loss to each node to optimize our CSE model, which can be represented by

$$\mathcal{L}_{cse} = \frac{1}{N} \sum_{p \in P} -\ln(y^*(p)), \quad (13)$$

where $N = S \times S$ represents the number of nodes in a grid, P denotes a set of all nodes, and y^* is the value of the positive merging direction. Our optimization strategy computes the loss according to the current CSE prediction. Intuitively, it intends to boost the positive candidates which are already strong, which reduces the ambiguity in labeling and speed-up the convergence.

³At most two merging directions will be labeled as positive.

4. Experiment

4.1. Experiment Details

The experiment is conducted on Tensorflow 1.5.0 [3]. We adopt Faster RCNN driven by ResNet-34 to initialize seed locations and corresponding grids in all experiments. The CSE is implemented and optimized in C++ and accelerated by CUDA. Following the existing training strategies for scene text detection [43, 27, 21], we pretrain our model on a combined dataset. The pretraining dataset consists of over 10k images from full set of ICDAR-17 MLT [1] and the training sets of MSRA-TD500 [37], Total-Text [5] and CTW-1500 [39]. After the pretraining, we fine-tune and evaluate our method on two curve text datasets Total-Text (with 1255 training images and 300 testing images) and CTW-1500 (with 1000 training images and 500 testing images). The model is trained on the combined dataset for 50k iteration and fine-tuned on the datasets to be evaluated. We adopt the Adam optimizer [16] to train the network. In the pretraining phase, the learning rate is fixed to 0.01 for the first 30k iterations and scaled down to 0.002 for the rest iterations. In the fine-tuning, the initial learning rate is set to 0.001 and decays exponentially 0.9 every 5000 iterations. All the experiment is conducted on Dell Alienware with Intel i7 processor, 64GB memory and two NVIDIA GTX 1080 Ti GPUs. The batch size is set to 1 for each of two GPUs in training and only one GPU is used for evaluation.

4.2. Flexibility and Robustness

In this experiment, we validate the flexibility and robustness of our CSE method qualitatively and quantitatively. In the qualitative experiment, we generate a set of sampling grids (in yellow) with different locations and sizes by randomly manipulating the ground-true boxes as shown in the first row of Fig. 6. We apply the CSE to the corresponding RoIs shown in the second row of Fig. 6, and visualize the extracted text regions by heat maps in the third row of Fig. 6. Fig. 6 (a) shows the flexibility of our method. Our CSE method can effectively retrieve the text region with different seed's locations. Even for a text object with large curvature and slim shape, our method can capture all the related sub-regions with high area precision. On the other hand, for a proposed region with many unexpected texts included or even dominated by another text instances (demonstrated in Fig. 6 (b) and (c)), our method only extracts associated object regions indicated by the seed. It indicates that our CSE is robust to the ambiguity caused unexpected objects and can produce satisfactory results even for a poor sampling grid generated by the previous object detector.

In addition to visually investigating our CSE, we quantitatively verify its flexibility and robustness by rescaling the size of a proposed sampling grid and

relocating the seed in a grid. The grid rescaling resizes the height and width of a proposed region proposal by a factor of $\delta_s \geq 1.0$. The seed relocation is applied to a sampling grid to change its seed to a new node which is still within the targeted object region but have δ_c deviation in Euclidean space from the original seed node. δ_c is normalized by the square root of the original grid area. We study the effects of rescaling and relocation separately by profiling the precisions, recalls, and F-scores on both Total-Text and CTW-1500, and the results are shown in Fig. 7. The performance variation respective to rescaling factor δ_s on two datasets is profiled in Fig. 7 (a) and (b). The performance is maintained at around 80% on Total-Text and 78% on CTW-1500 for δ_s ranging from 1.0 to 1.5. It slightly drops when δ_s is larger than 1.5 and the F-scores remain above 77% and 73% respectively. As for seed relocation, the according performance variations on Total-Text and CTW-1500 are illustrated in Fig. 7 (c) and (d). The performance is not greatly affected by the seed relocation. The F-scores remains 80% and 78% and decrease by only 3% and 4% when δ_c is changed from 0.0 to 1.0. In conclusion, our CSE is robust to randomly initialized seed locations and distorted sampling grids.

4.3. Comparing with Mask RCNN

The baseline Mask RCNN method is implemented based on the method proposed in [27]. We remove the text recognition branch and only keep the detection and segmentation branches. For a fair comparison, two methods share the same text detector which is based on Faster RCNN architecture. The quantity results are reported in Tab 4.3. Our method is overall better than the baseline method. The F-scores of the baseline method on both datasets are 67.5% to 67.8%, respectively. In comparison, our CSE performs much better than the Mask RCNN based method by over 10% with F-score of 80.2% on Total-Text and 77.4% on CTW-1500.

To explore the cause of performance gain, we visually compare the output confidence map produced by two methods on CTW-1500, which is shown in Fig. 8. Fig. 8 (a) demonstrates the failed examples produced by the baseline method. In these cases, the segmentation is distorted by the adjacent text instance. Parts of the unexpected text instances included in a box cause high activation in a confidence map and corrupt the boundary prediction. In contrast, our CSE is extremely robust in this scenario, since the contextual information captured by CSE helps to eliminate the ambiguity caused by unexpected objects. Moreover, the condition modeling allows our CSE to retrieve long curve text lines with high precision, which is flexible and promising in real applications.

Figure 6. Robustness and Flexibility Analysis: In the first row of each case, the sampling regions are represented by the bounding boxes in yellow and the seed locations are labeled by the cross dash lines. The second row shows the zoom-in of RoIs. The corresponding heatmaps of associate regions are shown in the third row. As shown in (a), our method is very flexible in the seed’s locations. (b) and (c) prove the robustness of our method to extract text from a density text region.

Figure 7. Performance v.s. rescaling factor δ_s on Total-Text (a) and CTW-1500 (b); Performance v.s. relocation factor δ_c on Total-Text (c) and CTW-1500 (d).

Datasets	Total-Text			CTW1500		
Model	P	R	F	P	R	F
MRCNN	69.2	65.8	67.5	65.1	70.8	67.8
CSE	81.4	79.1	80.2	78.7	76.1	77.4

Table 1. Performance Comparison between Mask RCNN based method and our CSE method.

4.4. Comparing with Polygon Regression

In this experiment, we compare our method with another baseline based on polygon regression proposed in [39]. The baseline is implemented based on the publicly available source code provided in [39], and is pretrained and fine-tuned with our training strategy. Similar to the previous experiment, the backbone and the RPN are shared among two methods, and the rest parts are implemented based on respective workflows. The results are shown in Tab 4.4. Our

Figure 8. (a) Examples output by baseline method; (b) Examples output by our CSE method. Compared with the baseline Mask RCNN method, our method shows impressing robustness to the ambiguity caused by adjacent texts and impressing selectivity to the targeted instance.

method outperforms the baseline method by 5% in terms of F-score on Total-Text, with the precision of 80.9%, recall of 80.3% and F-score of 80.6%. On CTW-1500, our CSE achieves F-score of 77.6 %, which is 4.4% better than the baseline.

We also investigate the causes of performance gain by visualizing the detection results of both methods. As demonstrated in Fig. 9, with the same RoI proposals shown in Fig. 9 (a), the polygon regression could be corrupted by the other text object which is occasionally included. As can be seen in Fig. 9 (b), the baseline model may consider all the texts included in a proposed region as a single object and regress the corresponding boundary. Although this problem can be mitigated by training a more accurate text detector to reduce the unexpected texts, the proposed bounding boxes inevitably cover additional texts due to text’s highly varying shapes and orientations. As

Datasets	Total-Text			CTW1500		
Methods	P	R	F	P	R	F
Poly-Reg	73.8	77.4	75.6	77.1	69.7	73.2
CSE	80.9	80.3	80.6	79.2	76.0	77.6

Table 2. Performance Comparison between Polygon Regression based method and our CSE method.

Figure 9. (a) The images with RoIs; (b) The text boundaries output by baseline method; (c) The text boundaries produced by our CSE method. The baseline method is easily affected by the unexpected texts included in the same boxes, while our method shows strong robustness to this situation.

shown in Fig. 9 (c), the conditional expansion mechanism only merges the sub-regions that are similar to the region indicated by the seed. By exploring the spatial dependency as well as the local information, our CSE method is much more robust than the polygon regression method and can produce more elaborated boundaries.

4.5. Peer Comparison

We compare our method with the recently proposed methods for curve text detection on curve text benchmarks, Total-Text and CTW-1500. The results are shown in Tab.4.5 and Tab.4.5. Our method creates a new state-of-art performance with the precision of 81.4%, recall of 79.1% and F-score of 80.2% on Total-Text. On CTW-1500 containing both curve texts and long text lines, our method also achieves the state-of-art performance with the precision of 81.1%, recall of 76.0% and F-score of 78.4%. The inference time is 0.42 ms per image and 0.38 ms per image on Total-Text and CTW-1500 respectively. The detection results are demonstrated in Fig. 10. It shows that our method can effectively handle curve texts with irregular shapes, highly varying sizes and arbitrary orientations.

Methods	P	R	F	time (s)
SegLink [32]	30.3	23.8	26.7	-
EAST [43]	50.0	36.2	42.0	-
Mask TextSpotter [27]	69.0	55.0	61.3	-
TextSnake et al. [26]	82.7	74.5	78.4	-
CSE	81.4	79.1	80.2	0.42

Table 3. Detection Performance on Total-Text.

Methods	P	R	F	time (s)
SegLink [32]	42.3	40.0	40.8	-
EAST [43]	78.7	49.1	60.4	-
DMPNet [22]	69.9	56.0	62.2	-
CTD [39]	74.3	65.2	69.5	-
CTD+TLOC [39]	77.4	69.8	73.4	-
TextSnake et al. [26]	67.9	85.3	75.6	-
CSE	81.1	76.0	78.4	0.38

Table 4. Detection Performance on CTW-1500.

Figure 10. Detection results on Total-Text and CTW-1500.

5. Conclusion

In this work, we analyze the deficiency of the existing curve text detection methods and improve the performance by developing a novel parameterized Conditional Spatial Expansion (CSE) mechanism. Our method shows strong robustness to the ambiguity caused by close texts with arbitrary shapes and orientations. It is flexible and can extract text regions in a controllable manner. Our CSE method outperforms the existing curve text detection methods.

Acknowledgement

G. Lin's participation was partly supported by the National Research Foundation Singapore under its AI Singapore Programme [AISG-RP-2018-003] and a MOE Tier-1 research grant [RG126/17 (S)]. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.

References

- [1] Icdar 2017 robust reading competition. <http://u-pat.org/ICDAR2017/index.php>. 6
- [2] Resnet-34. <https://www.kaggle.com/pytorch/resnet34>.
- [3] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. 6
- [4] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [5] Chee Kheng Ch'ng and Chee Seng Chan. Total-text: A comprehensive dataset for scene text detection and recognition. In *Document Analysis and Recognition (ICDAR), 2017 14th IAPR International Conference on*, volume 1, pages 935–942. IEEE, 2017. 2, 6
- [6] Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. Pixellink: Detecting scene text via instance segmentation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018. 2
- [7] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Context contrasted feature and gated multi-scale aggregation for scene segmentation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 2
- [8] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Semantic correlation promoted shape-variant context for segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [9] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 2
- [10] Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. An end-to-end textspotter with explicit alignment and attention. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5020–5029, 2018. 2
- [11] Wenhao He, Xu-Yao Zhang, Fei Yin, and Cheng-Lin Liu. Deep direct regression for multi-oriented scene text detection. *arXiv preprint arXiv:1703.08289*, 2017. 2
- [12] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 4
- [13] Weilin Huang, Yu Qiao, and Xiaoou Tang. Robust scene text detection with convolution neural network induced msr trees. In *European Conference on Computer Vision*, pages 497–511. Springer, 2014. 2
- [14] Yingying Jiang, Xiangyu Zhu, Xiaobing Wang, Shuli Yang, Wei Li, Hua Wang, Pei Fu, and Zhenbo Luo. R2cnn: rotational region cnn for orientation robust scene text detection. *arXiv preprint arXiv:1706.09579*, 2017. 1, 2
- [15] 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 *Document Analysis and Recognition (ICDAR), 2015 13th International Conference on*, pages 1156–1160. IEEE, 2015.
- [16] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [18] Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. Textboxes: A fast text detector with a single deep neural network. In *AAAI*, pages 4161–4167, 2017. 2
- [19] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, volume 1, page 4, 2017.
- [20] 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
- [21] Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. Fots: Fast oriented text spotting with a unified network. *arXiv preprint arXiv:1801.01671*, 2018. 2, 6
- [22] Yuliang Liu and Lianwen Jin. Deep matching prior network: Toward tighter multi-oriented text detection. 8
- [23] Z. Liu, G. Lin, W. L. Goh, F. Liu, C. Shen, and X. Yang. Correlation Propagation Networks for Scene Text Detection. *ArXiv e-prints*, Sept. 2018. 1, 2
- [24] Zichuan Liu, Guosheng Lin, Sheng Yang, Jiashi Feng, Weisi Lin, and Wang Ling Goh. Learning markov clustering networks for scene text detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1, 2
- [25] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [26] S. Long, J. Ruan, W. Zhang, X. He, W. Wu, and C. Yao. TextSnake: A Flexible Representation for Detecting Text of Arbitrary Shapes. *ArXiv e-prints*, July 2018. 3, 8
- [27] P. Lyu, M. Liao, C. Yao, W. Wu, and X. Bai. Mask TextSpotter: An End-to-End Trainable Neural Network for Spotting Text with Arbitrary Shapes. *ArXiv e-prints*, July 2018. 1, 3, 6, 8
- [28] Pengyuan Lyu, Cong Yao, Wenhao Wu, Shuicheng Yan, and Xiang Bai. Multi-oriented scene text detection via corner localization and region segmentation. *arXiv preprint arXiv:1802.08948*, 2018. 1
- [29] Jianqi Ma, Weiyuan Shao, Hao Ye, Li Wang, Hong Wang, Yingbin Zheng, and Xiangyang Xue. Arbitrary-oriented scene text detection via rotation proposals. *IEEE Transactions on Multimedia*, 2018. 1, 2
- [30] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object

- detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 779–788, 2016. 2
- [31] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. 2, 5
- [32] Baoguang Shi, Xiang Bai, and Serge Belongie. Detecting oriented text in natural images by linking segments. *arXiv preprint arXiv:1703.06520*, 2017. 1, 8
- [33] Abhinav Shrivastava, Abhinav Gupta, and Ross Girshick. Training region-based object detectors with online hard example mining. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 761–769, 2016.
- [34] Zhi Tian, Weilin Huang, Tong He, Pan He, and Yu Qiao. Detecting text in natural image with connectionist text proposal network. In *European Conference on Computer Vision*, pages 56–72. Springer, 2016. 1
- [35] Yue Wu and Prem Natarajan. Self-organized text detection with minimal post-processing via border learning. In *Proc. ICCV*, 2017. 2
- [36] Chuhui Xue, Shijian Lu, and Fangneng Zhan. Accurate scene text detection through border semantics awareness and bootstrapping. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [37] Cong Yao, Xiang Bai, Wenyu Liu, Yi Ma, and Zhuowen Tu. Detecting texts of arbitrary orientations in natural images. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 1083–1090. IEEE, 2012. 6
- [38] Cong Yao, Xiang Bai, Nong Sang, Xinyu Zhou, Shuchang Zhou, and Zhimin Cao. Scene text detection via holistic, multi-channel prediction. *arXiv preprint arXiv:1606.09002*, 2016. 2
- [39] Liu Yuliang, Jin Lianwen, Zhang Shuaitao, and Zhang Sheng. Detecting curve text in the wild: New dataset and new solution. *arXiv preprint arXiv:1712.02170*, 2017. 1, 2, 6, 7, 8
- [40] Fangneng Zhan and Shijian Lu. Esir: End-to-end scene text recognition via iterative image rectification. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1
- [41] Fangneng Zhan, Shijian Lu, and Chuhui Xue. Verisimilar image synthesis for accurate detection and recognition of texts in scenes. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [42] Zheng Zhang, Chengquan Zhang, Wei Shen, Cong Yao, Wenyu Liu, and Xiang Bai. Multi-oriented text detection with fully convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4159–4167, 2016. 2
- [43] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: an efficient and accurate scene text detector. *arXiv preprint arXiv:1704.03155*, 2017. 2, 6, 8