GyF

This ACCV 2020 paper, provided here by the Computer Vision Foundation, is the author-created version. The content of this paper is identical to the content of the officially published ACCV 2020 LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/accv

Synthetic-to-Real Unsupervised Domain Adaptation for Scene Text Detection in the Wild

Weijia Wu¹, Ning Lu², Enze Xie³, Yuxing Wang¹, Wenwen Yu⁴, Cheng Yang¹, and Hong Zhou^{1 \boxtimes}

 ¹ Zhejiang University, Key Laboratory for Biomedical Engineering of Ministry, HangZhou, China. zhouhong_zju@126.com
 ² Tencent Cloud Product Department
 ³ The University of Hong Kong
 ⁴ Xuzhou Medical University

Abstract. Deep learning-based scene text detection can achieve preferable performance, powered with sufficient labeled training data. However, manual labeling is time consuming and laborious. At the extreme, the corresponding annotated data are unavailable. Exploiting synthetic data is a very promising solution except for domain distribution mismatches between synthetic datasets and real datasets. To address the severe domain distribution mismatch, we propose a synthetic-to-real domain adaptation method for scene text detection, which transfers knowledge from synthetic data (source domain) to real data (target domain). In this paper, a text self-training (TST) method and adversarial text instance alignment (ATA) for domain adaptive scene text detection are introduced. ATA helps the network learn domain-invariant features by training a domain classifier in an adversarial manner. TST diminishes the adverse effects of false positives (FPs) and false negatives (FNs) from inaccurate pseudo-labels. Two components have positive effects on improving the performance of scene text detectors when adapting from synthetic-to-real scenes. We evaluate the proposed method by transferring from SynthText, VISD to ICDAR2015, ICDAR2013. The results demonstrate the effectiveness of the proposed method with up to 10%improvement, which has important exploration significance for domain adaptive scene text detection. Code is available at https://github.com/ weijiawu/SyntoReal_STD.

1 Introduction

Scene text detection and recognition[1–3] has received increasing attention due to its numerous applications in computer vision. Additionally, scene text detection [4–8] has achieved great success in the last few decades. However, these detection methods require manually labeling large quantities of training data, which is very expensive and time consuming. Whereas several public benchmarks [9–13] have already existed, they only covered a very limited range of scenarios. In the real world, a specific application task usually requires the collection and annotation of a new training dataset, and it is difficult, even impossible, to collect enough labeled data. Therefore, the expensive cost of labeling has



Fig. 1. Examples of different datasets. The first row are from real ICDAR2013[9], IC-DAR2015[10], and ICDAR2017 MLT[11], respectively. The second row is from Virtual SynthText[14], VISD[15], and UnrealText[16]. There remains a considerable domain gap between synthetic data and real data.

become a major problem for text detection applications based on deep learning methods.

With the great development of computer graphics, an alternative way is to utilize synthetic data, which is largely available from the virtual world, and the ground truth can be freely and automatically generated. SynthText [14] first provides a virtual scene text dataset and automatically generates synthetic images with word-level and character-level annotations. Zhan etal. [15] equipped text synthesis with selective semantic segmentation to produce more realistic samples. UnrealText [16] provides realistic virtual scene text images via a 3D graphics engine, which provides realistic appearance by rendering scene and text as a whole. Although synthetic data offer the possibility of substituting for real images in training scene text detectors, many previous works have also shown that training with only synthetic data degrades the performance on real data due to a phenomenon known as "domain shift". As shown in Fig. 1, unlike common objects, text has more diversity of shapes, colours, fonts, sizes, and orientations in real-world scenarios, which causes a large domain gap between synthetic data and real data. Therefore, the performance of the model degrades significantly when applying model learning only from synthetic data to real data.

To tackle the domain shift, we propose a synthetic-to-real domain adaptation approach for scene text detection, which aims to efficiently improve the model performance on real data by using synthetic data and unlabeled real data. Inspired by [17] and [18], a text self-training(TST) method and an adversarial text instance alignment(ATA) are proposed in our paper to reduce the domain shift. Self-training has achieved excellent results for domain adaptive object detection [19, 20] and semantic segmentation [21]. However, scene text detection tasks with more diverse situations and complex backgrounds have not been explored in this direction to the best of our knowledge. To better apply self-training to scene text detection, TST is used to suppress the adverse impact of both false positives and false negatives that occur in pseudo-labels. In addition, we first utilize adversarial learning [18] help the model to learn discriminative features of text. The contributions of our paper are as follows:

- We introduce text self-training (TST) to improve the performance of domain adaptive scene text detection by minimizing the adverse effects of inaccurate pseudo-labels.
- We propose adversarial text instance alignment (ATA) to help the model learn domain-invariant features, which enhance the generalization ability of the model.
- We first introduce a synthetic-to-real domain adaptation method for scene text detection, which transfers knowledge from the synthetic data (source domain) to real data (target domain).

The proposed method is evaluated by extensive experiments for the scene text detection transfer task(*e.g.*, SynthText [14] \rightarrow ICDAR2015 [10]). The experimental results demonstrate the effectiveness of the proposed approach for addressing the domain shift of scene text detection, which has important exploration significance for domain adaptive scene text detection.

2 Related Work

2.1 Scene Text Detection

Before the era of deep learning, SWT [22] and MSER [23] were two representative algorithms for conventional text detection methods. SWT obtains information about the text stroke efficiently, and MSER draws intensity stable regions as text candidates. Based on convolutional neural network knowledge, scene text detection [24, 25] has made great progress. EAST [8] performs very dense predictions that are processed using locality-aware NMS. PixelLink [26] detects text instances by linking neighbouring pixels. PSENet [27] proposed a progressive scale algorithm to gradually expand the predefined kernels for scene text detection. In addition to the above methods based on strongly supervised learning, some weakly/semi-supervised methods are proposed to reduce the expansive cost of annotation. WeText [28] trains a text detection model on a small amount of character-level annotated text images, followed by boosting the performance with a much larger amount of weakly annotated images at word line level. WordSup [29] trains a character detector by exploiting word annotations in rich, large-scale real scene text datasets. [30] utilizes the network pretrained on synthetic data with full masks to enhance the coarse masks in a real image.

2.2 Domain Adaptation

Domain adaptation reduces the domain gap between training and testing data. Prior works [31] estimated the domain gap and minimized it. Recent methods

use more effective methods to reduce the domain gap, such as incorporating a domain classifier with gradient reversal [18]. [32] addressed the domain shift by training domain discriminators on the image level and instance level. [17] introduced a weak self-training to diminish the adverse effects of inaccurate pseudo-labels, and designed adversarial background score regularization to extract discriminative features. For scene text, the domain adaptation method [33] converts a source-domain image into multiple images of different spatial views as in the target domain. Handwriting recognition [34] proposes AFDM to elastically warp extracted features in a scalable manner.

2.3 Self-Training

Prior works used self-training [35, 36] to compensate for the lack of categorical information. [37] bridged the gap between the source and target domains by adding both the target features and instances in which the current algorithm is the most confident. [38] used three networks asymmetrically, where two networks were used to label unlabeled target samples and one network was trained to obtain discriminative representations. Other works [39, 40, 21] also showed the effectiveness of self-training for domain adaptation. However, text detection still requires further exploration in the self-training method due to a lack of previous work.

3 Proposed Method

In this section, the problems caused by domain shifts are analysed. Furthermore, we introduce the principle of TST and ATA, and how to use them for domain adaptation. To evaluate our method, EAST [8] is adopted as the baseline.

3.1 Problem and Analysis

Although synthetic scene text data can be automatically generated with diversified appearance and accurate ground truth annotations, the model trained with only synthetic data cannot be directly applied to real scenes since there exists a significant domain shift between synthetic datasets and real datasets.

Viewing the problem from a probabilistic perspective is clearer. We refer to the synthetic data domain as the source domain and the real data domain as the target domain. The scene text detection problem can be viewed as learning the posterior P(B|I), where I refers to the image features and B is the predicted bounding-box of text instances. Using the Bayes formula, the posterior P(B|I) can be decomposed as:

$$P(B|I) = \frac{P(I|B) * P(B)}{P(I)} = \frac{P(I|B)}{P(I)} * P(B).$$
(1)

We make the covariate shift assumption in this task that the priori probability P(B) is the same for the two domains. P(I|B) refers to the conditional probability of I, which means that the likelihood of learning true features given that



Fig. 2. The network architecture with the corresponding optimization object. θ represents the parameters of EAST. A domain classifier (green) is added after the feature extractor via a gradient reversal layer that multiplies the gradient by a certain negative constant during the backpropagation-based training. L_{task} refers to the original detection loss of EAST, and L_d is the loss of domain classifier.

the predicted result is true. We also consider that P(I|B) is the same for both domains. Therefore, the difference in posterior probability is caused by the priori probability P(I). In other words, to detect text instances, the difference in detection results is caused by domain change features. To improve the generalization ability, the model should learn more domain-invariant features, keeping the same P(I) regardless of which domain the input image belongs.

In the EAST [8] model, the image feature P(I) refers to the features output from the backbone. Therefore, the feature map should be aligned between the source domain and the target domain (*i.e.*, $P_s(I) = P_t(I)$). To achieve this goal, ATA is proposed to align the features, with more details in the next subsection.

3.2 Adversarial Text Instance Alignment

Motivated by [18], ATA is adopted to help the network learn domain-invariant features. In the EAST model, the image features P(I) refer to the feature map outputs of the backbone (*i.e.*, 384, 1/32 features in Fig. 2). To align the features P(I) between the source domain and target domain, a domain classifier is used to confuse the feature domain.

In particular, the domain classifier is trained for each input image and predicts the domain label to which the image belongs. We assume that the model works with input samples $x \in X$, where X is the some input space. y_i denotes the domain label of the *i*-th training image, with $y_i = 0$ for the source domain and $y_i = 1$ for the target domain. $p_i(x)$ is the output of the domain classifier, and we use cross entropy as the loss function:

$$L_d = -\sum_i (y_i \times ln^{p_i(x)} + (1 - y_i) \times ln^{1 - p_i(x)}).$$
(2)

To learn domain-invariant features, we optimize the parameters in an adversarial way. The parameters of the domain classifier are optimized by minimizing the above domain classification loss, and the parameters of the base network are optimized by maximizing this loss. For more detail, the gradient reverse layer (GRL) [18] is added between the backbone of EAST and the domain classifier, and the sign of the gradient is reversed when passing through the GRL layer.

As shown in Fig. 2, both the feature pyramid network (FPN) and the backbone minimize the original loss L_{task} of EAST at the training phase. L_{task} specifically denotes the score map loss and geometries loss in EAST [8]. L_{task}^t refers to training with the pseudo-label in the target domain, and L_{task}^s denotes training with the source domain. Thus, different training objectives for various parameter spaces:

$$\begin{cases} L_f = \min(L_{task}^t(\theta_f|x^t) + L_{task}^s(\theta_f|x^s) - \lambda L_d(\theta|(x^s, x^t)))) & \theta_f \in F, \\ L_d = \min(L_d(\theta_d|(x^s, x^t))) & \theta_d \in C, \\ L_h = \min(L_{task}^t(\theta_h|x^t) + L_{task}^s(\theta_h|x^s)) & \theta_h \in D, \end{cases}$$
(3)

where F, C, D are the parameter spaces of the backbone, the domain classifier and the FPN. The overall training objective is as follows:

$$L = L_f + L_h + \lambda L_d \,, \tag{4}$$

where λ is the tradeoff parameter, we set it to 0.2 in all experiments. Through optimizing the loss, the network can learn more text domain-invariant features, transforming better from synthetic data to real data.

3.3 Text Self-Training

Previous works [21, 41] have shown the effectiveness of self-training. However, two major problems for self-training still need to be explored further: false positives(FP) and false negatives(FN) occurred in pseudo-label. Incorrect pseudo-labels will cause very serious negative effects to our networks. To overcome such problems, TST is designed to minimize the adverse effects of FP and FN.

Reducing False Negatives. Inspired by [17], a weak supervision way is utilized to minimize the effects of false negatives. The original score map loss in the EAST [8] is

$$L_s = -\sum_{i \in Pos} \beta Y^* log \widehat{Y} - \sum_{i \in Neg} (1-\beta)(1-Y^*)(1-\widehat{Y}), \qquad (5)$$

where $\hat{Y} = F_s$ is the prediction of the score map, and Y^* is the ground truth. While the network is optimized by backpropagation learning the loss of background (*i.e.*, negative examples), FP occurring in pseudo-labels misleads the network. We assume that FPs are mainly selected by hard negative mining,



Fig. 3. Up: The framework of proposed text self-training. We utilize SWS to filter the positive samples in pseudo-label for minimizing false positives, and select a third of negative samples with low confidence as the final negative samples to minimize false negatives. Down: We present sample space representation for pseudo-label. (a): False negatives are effectively filtered out by weak training. (b): False positives are filtered out by the standard deviation(σ) of the stroke width and SWS.

such as blurred text and unusual fonts similar to the background. To reduce the adverse effects of FP, we ignore some background examples that have the potential to be foregrounds with a confidence score.

Negative examples for EAST are a pixel map, a pixel is more likely to be considered a negative pixel while the corresponding confidence score higher. Thus, we choose a part of the negative sample pixels(*e.g.*, Neg/3) that have the lowest confidence score as the final negative examples, which is denoted as \widehat{Neg} in Fig. 3(red line). The corresponding mathematical expression is $\widehat{Neg} = \eta Neg$, where η is set to 1/3 in all experiments. For those pixels that have a high confidence score, the network does not optimize this part loss. Finally, the modified loss function is defined as

$$L_{sw} = -\sum_{i \in Pos} \beta Y_i^* log \widehat{Y} - \sum_{i \in \widehat{Neg}} (1-\beta)(1-Y_i^*)(1-\widehat{Y}).$$
(6)

Reducing False Positives. Corresponding to false negatives, false positives also cause serious interference to the network. Some patterns and designs in natural scenes are extremely easy to identify as text, which leads to inaccurate pseudo-labels. Replacing Supporting Region-based Reliable Score(SRRS) in [17], we propose a more reasonable Stroke Width Score(SWS) that utilizes the Stroke Width Transform(SWT) [22] to evaluate the predicted boxes of text instances. On the one hand, SRRS is not applicable to EAST based on segmentation. SRRS

in [17] is define as:

$$SRRS(r^*) = \frac{1}{N_s} \sum_{i=1}^{N_s} IoU(r_i, r^*) \cdot P(c^* | r_i)$$
(7)

The EAST is a segmentation-based method without FPN, the text instances with small area have less supporting boxes $(i.e.,r_i)$ than that with big area, which leads to the extremely unbalanced supporting boxes number $(i.e.,N_s)$. On the other hand, SWT is more reasonable for eliminating non-text regions, and similar previous works [42, 13] have shown its effectiveness.

SWT is a local image operator that computes the most likely stroke width for each pixel. The output of the SWT is a n * m matrix where each element contains the width of the stroke associated with the pixel. Specifically, each pixel on the boundary of the stroke is connected with the opposite side of the stroke in the direction of the gradient, and the width of the connecting line is the width of the stroke for the pixel. SWS assesses the predicted boxes by utilizing the information of the corresponding stroke width, and eliminates part of the non-text regions, as shown in Fig. 3(blue line).

For a typical text region, the variance in stroke width is low since text tends to maintain a fixed stroke width. We denote the set of stroke widths in the v_{th} predicted box as W_n^v and the stroke width of the u_{th} pixel as $w_u^v \in W_n^v$. The standard deviation is as follows:

$$\sigma_v = \sqrt{\frac{1}{N} \sum_{u=1}^{N} (w_u^v - \mu^v)^2},$$
(8)

where μ^{v} is the mean stroke width in the v_{th} predicted box. Therefore, each box has a standard deviation(σ) about the stroke width, and we choose reliable boxes with an upper threshold(ϵ_1). In addition, we further filter boxes by SWS:

$$SWS_v = \frac{w_v}{\sigma_v^2} \,, \tag{9}$$

 w_v is the most common stroke width value for the v_{th} predicted box. By thresholding the score with a lower threshold ϵ_2 , the boxes are further selected. Fig. 3 (b) shows that part of the FP is filtered out by SWS and σ .

4 Experiments

The proposed method is evaluated by transferring a scene text detector from synthetic datasets to real datasets. We adopt several pure synthetic data and real scene data (*i.e.*, SynthText [14] and ICDAR2015 [10]).

4.1 Dataset and Experimental Settings

Pure Synthetic Datasets. SynthText [14] is a large-scale dataset that contains about 800K synthetic images. These images are created by blending natural images with text rendered with random fonts, sizes, colours, and orientations, thus

Method	Source \rightarrow Target(UL)	Annotation	Detection Evaluation/%			
		Annotation	Precision	Recall	F-score	
PAN [43]	$SynthText \rightarrow ICDAR2015$	Word	0.659	0.469	0.548	
EAST [8]*	$SynthText \rightarrow ICDAR2015$	Word	-	-	0.580	
EAST $[8]^{\dagger}$	$SynthText \rightarrow ICDAR2015$	Word	0.721	0.521	0.605	
CCN [44]	$SynthText \rightarrow ICDAR2015$	Character	-	-	0.651	
EAST+Ours	$SynthText \rightarrow ICDAR2015$	Word	0.690	0.670	0.680	
EAST [8]*	$VISD \rightarrow ICDAR2015$	Word	-	-	0.643	
EAST [8] [†]	$VISD \rightarrow ICDAR2015$	Word	0.640	0.652	0.645	
EAST+Ours	$VISD \rightarrow ICDAR2015$	Word	0.748	0.727	0.738	

Table 1. The performance of different models on Syn2Real scene text detection dataset for SynthText/VISD \rightarrow ICDAR2015 transfers. UL refers to the unlabeled data. * denotes the performance reported in UnrealText [16]. † refers to our testing performance.

these images are quite realistic. *Verisimilar Image Synthesis Dataset(VISD)* [15] contains 10 k images synthesized with 10 k background images. Thus, there are no repeated background images for this dataset.

Real Datasets. *ICDAR2015* [10] is a multi-oriented text detection dataset for English text that includes 1,000 training images and 500 testing images. Scene text images in this dataset were taken by Google Glasses without taking care of positioning, image quality, and viewpoint. *ICDAR2013* [9] was released during the ICDAR 2013 Robust Reading Competition for focused scene text detection, consisting of high-resolution images, 229 for training and 233 for testing, containing texts in English.

Implementation Details. In all experiments, we used EAST [8] as a base network. Following the original paper, inputs were resized to 512×512 , and we applied all augmentations used in the original paper. The network was trained with a batch input composed of 12 images, 6 images from the source domain, and the other 6 images from the target domain. The Adam optimizer was adopted as our learning rate scheme. All of the experiments used the same training strategy: (1) pretraining the network for 80 k iterations with ATA to learn domaininvariant features and (2) the pretrained model is used to generate corresponding pseudo-label(*i.e.*,pseudo-bounding box label and negative sample map) for each image in the target domain, then fine-tuning the pretrained model with generated pseudo-labels. In the process of generating pseudo-labels, we set ϵ_1 and ϵ_2 to 3.0 and 0.30 for stroke width elimination parameters. All of the experiments were conducted on a regular workstation (CPU: Intel(R) Core(TM) i7-7800X CPU @ 3.50 GHz; GPU: GTX 2080Ti).

4.2 Performance Comparison and Analysis

Synthetic \rightarrow ICDAR2015 Transfer. Table 1 summarizes the performance comparisons for synthetic \rightarrow ICDAR2015 transfer task. The EAST model as the

Mathad	Source > Target(III)	Appotation	Detection Evaluation/%			
Method	$300000 \rightarrow 100000000$	Annotation	Precision	Recall	F-score	
EAST [8]*	$SynthText \rightarrow ICDAR2013$	Word	-	-	0.677	
EAST $[8]^{\dagger}$	$SynthText \rightarrow ICDAR2013$	Word	0.669	0.674	0.671	
EAST+Ours	$SynthText \rightarrow ICDAR2013$	Word	0.805	0.765	0.784	
EAST [8]*	$VISD \rightarrow ICDAR2013$	Word	-	-	0.748	
EAST [8] [†]	$VISD \rightarrow ICDAR2013$	Word	0.783	0.705	0.742	
EAST+Ours	$VISD \rightarrow ICDAR2013$	Word	0.830	0.781	0.805	

Table 2. The performance of different models on Syn2Real scene text detection dataset for SynthText/VISD \rightarrow ICDAR2013 transfers. UL refers to the unlabeled data. * denotes the performance reported in UnrealText [16]. † refers to our testing performance.

baseline training with source-only had an unsatisfactory F-score(60.5% using SynthText and 64.5% using VISD), which can be regarded as a lower bound without adaptation. By combining with the proposed method, the F-score achieved a 68.0% and 73.8% respectively, making 7.5% and 9.3% absolute improvements over the baseline. GCN [44] based on character annotation led to a performance improvement over that based on word annotation. However, the performances of GCN were still lower than our method, which utilizes self-training and adversarial learning. The experiment indicates the efficient performance of the proposed method in alleviating the domain discrepancy over the source and target data.

Synthetic \rightarrow ICDAR2013 Transfer. To further verify the effectiveness of our proposed method, we conducted experiments by using ICDAR2013 as the target domain for the synthetic \rightarrow real scene text detection transfer task. The experimental results are reported in Table 2. Specifically, for the SynthText \rightarrow ICDAR2013 transfer task, compared with the baseline EAST training with source-only, we achieved an 11.3% performance improvement. Similar to synthetic \rightarrow ICDAR2015 transfer experiment, VISD was also used as the source domain in the comparison experiment. After using ATA and TST, the proposed method achieved a 6.3% performance improvement over the baseline EAST, which exhibits the effectiveness of the method for reducing the domain shift. Note that for fair comparison, except for adding ATA and TST, the base network and experimental settings of the proposed method were the same as the baseline EAST.

ICDAR2013 \rightarrow ICDAR2015 Transfer. Table 3 shows the performance for ICDAR2013 \rightarrow ICDAR2015 Transfer task. The annotations of ICDAR2013 are rectangular boxes while that of ICDAR2015 are rotated boxes, which limits the transfer performance. However, comparing with the baseline EAST training with source-only, we achieved an 7.6% performance improvement.

4.3 Ablation Study

Component Analysis. To verify the effectiveness of the proposed method, we conducted ablation experiments for Syn2Real transfer task on four datasets:

Table 3. Ablation study for the proposed Syn2Real scene text detection transfer. 'Baseline' denotes training only with labeled data in the source domain. \blacktriangle denotes the increase in the F-score compared with the baseline training with source-only. UL refers to the unlabeled data. 'F-target' denotes pretrain in source domain and fine-tuning with original pseudo-bounding box in target domain.

Method	TST		Source Termet(III)	Detection	Immer		
Method	151	AIA	$Source \rightarrow Larget(OL)$	Precision	Recall	F-score	improv.
Baseline				0.721	0.521	0.605	-
F-target	1			0.666	0.535	0.594	-
	\checkmark		$SynthText \rightarrow ICDAR2015$	0.693	0.605	0.646	▲4.1%
Ours		 ✓ 		0.682	0.610	0.640	▲3.5%
	\checkmark	 ✓ 		0.690	0.670	0.680	▲7.5%
Baseline				0.640	0.652	0.645	-
	\checkmark		VISD JCDAR2015	0.702	0.688	0.695	▲5.0%
Ours		 ✓ 	VISD-JICDAI(2013	0.713 0.670 0.691	0.691	▲4.6%	
	\checkmark	\checkmark		0.748	0.727	0.738	▲9.3%
Baseline				0.669	0.674	0.671	-
	√		Sumth Tourt VICD A D 2012	0.715	0.707	0.711	▲4.0%
Ours		\checkmark	Synth1ext→ICDAR2013	0.736	0.721	0.729	▲5.8%
	\checkmark	 ✓ 		0.805	0.765	0.784	▲11.3%
Baseline				0.783	0.705	0.742	-
	√			0.794	0.720	0.755	▲ 1.3%
Ours		 ✓ 	VISD-ICDAR2013	0.802	0.751	0.776	▲3.4%
	\checkmark	\checkmark		0.830	0.781	0.805	▲6.3%
Baseline				0.513	0.398	0.448	-
	√			0.546	0.459	0.505	▲5.7%
Ours		 ✓ 	$ 1 \cup DAR13 \rightarrow 1 \cup DAR2013 $	0.560	0.441	0.493	▲4.5%
	\checkmark	 ✓ 		0.563	0.490	0.524	▲7.6%

SynthText, VISD, ICDAR2015, and ICDAR2013. Table 3 shows the experimental results. For the SynthText \rightarrow ICDAR2015 transfer task, the F-scores increased by 4.1% and 3.5% respectively, after using the TST and ATA. In addition, our method produced a higher recall rate of up to eight percent than the baseline, which shows the effectiveness of this approach on improving the robustness of the model. By combining both components, the F-score of the proposed method achieved a 68.0%, a 7.5% absolute improvement over the baseline. The VISD \rightarrow ICDAR2015 transfer task exhibited better performance since VISD has a more realistic synthesis effect. In particular, the F-score using our method reached 73.8%, making the absolute improvement over the corresponding baseline 9.3%. For SynthText/VISD \rightarrow ICDAR2015 transfers, the improved performances are also significant. We achieved a 11.3% performance improvement using SynthText and 6.3% performance improvement using VISD.

Parameter Sensitivity on TST. To explore the influence of threshold parameters (*i.e.*, ϵ_1 and ϵ_2) on SWS, we conducted several sets of comparative experiments, and the results shown are in Table 4. Threshold parameter ϵ_1 was utilized to filter the predicted box since we considered the standard deviation of the stroke width in the text region close to zero in an ideal situation. The

Table 4. Model results for different values of ϵ_1 and ϵ_2 in Text Self-training(TST).

ϵ_2	ϵ_1	Detection Evaluation/%				62	Detection Evaluation/%		
		Precision	Recall	F-score	ε1	ϵ_2	Precision	Recall	F-score
-	-	0.597	0.561	0.580		-	0.597	0.561	0.580
	2	0.636	0.543	0.581		0.20	0.621	0.556	0.586
	3	0.634	0.563	0.596	-	0.30	0.623	0.554	0.586
	4	0.612	0.565	0.588		0.40	0.645	0.550	0.594



(a) SynthText-to-ICDAR2013 Transfer (b) SynthText-to-ICDAR2015 Transfer

Fig. 4. Examples of detection results for different models. The first row is the results of the baseline training with only source domain. The second row is the results of using the proposed method.

network trained with $\epsilon_1 = 3$ showed better performance than the others, and the results were not sensitive to the parameters. Similar to ϵ_1 , three different values 0.2, 0.3, 0.4 were adopted to verify the parameter sensitivity of ϵ_2 , and the result shows that the value(0.3) of ϵ_2 was reasonable.

Qualitative Analysis. Fig. 4 shows four examples of detection results for synthetic-to-real transfer tasks. The exemplars show that the proposed method improves the robustness of the model.

5 Conclusions

In this paper, we first introduced a synthetic-to-real domain adaptation method for scene text detection, which transfers knowledge from synthetic data to real data. The proposed TST effectively minimizes the adverse effects of FNs and FPs for pseudo-labels, and the ATA helps the network to learn domain-invariant features in an adversarial way. We evaluated the proposed method on several common synthetic and real datasets. The experiments showed that our approach makes a great improvement for synthetic-to-real transfer text detection task.

Acknowledgments. This work is supported by the National Key Research and Development Project (2019YFC0118202), the National Natural Science Foundation of China Grant No. 61803332 and No. LQ18E050001.

References

- He, W., Zhang, X.Y., Yin, F., Liu, C.L.: Deep direct regression for multi-oriented scene text detection. In: ICCV. (2017) 745–753
- Wang, W., Liu, X., Ji, X., Xie, E., Liang, D., Yang, Z., Lu, T., Shen, C., Luo, P.: Ae textspotter: Learning visual and linguistic representation for ambiguous text spotting. arXiv preprint arXiv:2008.00714 (2020)
- 3. Wang, W., Xie, E., Sun, P., Wang, W., Tian, L., Shen: Textsr: Content-aware text super-resolution guided by recognition. arXiv preprint arXiv:1909.07113 (2019)
- Liu, Z., Lin, G., Yang, S., Liu, F., Lin, W., Goh, W.L.: Towards robust curve text detection with conditional spatial expansion. In: CVPR. (2019) 7269–7278
- Long, S., Ruan, J., Zhang, W., He, X., Wu, W., Yao, C.: Textsnake: A flexible representation for detecting text of arbitrary shapes. In: ECCV. (2018) 20–36
- Xie, E., Zang, Y., Shao, S., Yu, G., Yao, C., Li, G.: Scene text detection with supervised pyramid context network. In: AAAI. (2019)
- Wang, W., Xie, E., Li, X., Hou, W., Lu, T., Yu, G., Shao, S.: Shape robust text detection with progressive scale expansion network. In: CVPR. (2019) 9336–9345
- Zhou, X., Yao, C., Wen, H., Wang, Y., Zhou, S., He, W., Liang, J.: East: an efficient and accurate scene text detector. In: CVPR. (2017) 5551–5560
- Karatzas, D., Shafait, F., Uchida, S., Iwamura, M., i Bigorda, L.G., Mestre, S.R., Mas, J., Mota, D.F., Almazan, J.A., De Las Heras, L.P.: Icdar 2013 robust reading competition. In: 2013 12th ICDAR, IEEE (2013) 1484–1493
- Karatzas, D., Gomez-Bigorda, L., Nicolaou, A., Ghosh, S., Bagdanov, A., Iwamura, M., Matas, J., Neumann, L., Chandrasekhar, V.R., Lu, S., et al.: Icdar 2015 competition on robust reading. In: 2015 13th ICDAR, IEEE (2015) 1156–1160
- Nayef, N., Yin, F., Bizid, I., Choi, H., Feng, Y., Karatzas, D., Luo, Z., Pal, U., Rigaud, C., Chazalon, J., et al.: Icdar2017 robust reading challenge on multilingual scene text detection and script identification-rrc-mlt. In: 2017 14th ICDAR. Volume 1., IEEE (2017) 1454–1459
- Yuan, T.L., Zhu, Z., Xu, K., Li, C.J., Hu, S.M.: Chinese text in the wild. arXiv preprint arXiv:1803.00085 (2018)
- 13. Yao, C., Bai, X., Liu, W., Ma, Y., Tu, Z.: Detecting texts of arbitrary orientations in natural images. In: CVPR, IEEE (2012)
- Gupta, A., Vedaldi, A., Zisserman, A.: Synthetic data for text localisation in natural images. In: CVPR. (2016) 2315–2324
- Zhan, F., Lu, S., Xue, C.: Verisimilar image synthesis for accurate detection and recognition of texts in scenes. In: ECCV. (2018) 249–266
- Long, S., Yao, C.: Unrealtext: Synthesizing realistic scene text images from the unreal world. arXiv preprint arXiv:2003.10608 (2020)
- Kim, S., Choi, J., Kim, T., Kim, C.: Self-training and adversarial background regularization for unsupervised domain adaptive one-stage object detection. In: ICCV. (2019) 6092–6101
- Ganin, Y., Lempitsky, V.: Unsupervised domain adaptation by backpropagation. arXiv preprint arXiv:1409.7495 (2014)
- Rosenberg, C., Hebert, M., Schneiderman, H.: Semi-supervised self-training of object detection models. WACV/MOTION 2 (2005)
- RoyChowdhury, A., Chakrabarty, P., Singh, A., Jin, S., Jiang, H., Cao, L., Learned-Miller, E.: Automatic adaptation of object detectors to new domains using selftraining. In: CVPR. (2019) 780–790

- 14 W. Wu et al.
- Zou, Y., Yu, Z., Vijaya Kumar, B., Wang, J.: Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In: ECCV. (2018) 289–305
- Epshtein, B., Ofek, E., Wexler, Y.: Detecting text in natural scenes with stroke width transform. In: CVPR, IEEE (2010) 2963–2970
- Neumann, L., Matas, J.: A method for text localization and recognition in realworld images. In: ACCV, Springer (2010)
- Xu, Y., Wang, Y., Zhou, W., Wang, Y., Yang, Z., Bai, X.: Textfield: Learning a deep direction field for irregular scene text detection. IEEE Transactions on Image Processing 28 (2019) 5566–5579
- Liao, M., Shi, B., Bai, X.: Textboxes++: A single-shot oriented scene text detector. IEEE transactions on image processing 27 (2018) 3676–3690
- Deng, D., Liu, H., Li, X., Cai, D.: Pixellink: Detecting scene text via instance segmentation. In: AAAI. (2018)
- Li, X., Wang, W., Hou, W., Liu, R.Z., Lu, T., Yang, J.: Shape robust text detection with progressive scale expansion network. arXiv preprint arXiv:1806.02559 (2018)
- Tian, S., Lu, S., Li, C.: Wetext: Scene text detection under weak supervision. In: ICCV. (2017) 1492–1500
- 29. Hu, H., Zhang, C., Luo, Y., Wang, Y., Han, J., Ding, E.: Wordsup: Exploiting word annotations for character based text detection. In: ICCV. (2017) 4940–4949
- 30. Wu, W., Xing, J., Yang, C., Wang, Y., Zhou, H.: Texts as lines: Text detection with weak supervision. Mathematical Problems in Engineering **2020** (2020)
- Gong, B., Shi, Y., Sha, F., Grauman, K.: Geodesic flow kernel for unsupervised domain adaptation. In: CVPR, IEEE (2012) 2066–2073
- Chen, Y., Li, W., Sakaridis, C., Dai, D., Van Gool, L.: Domain adaptive faster r-cnn for object detection in the wild. In: CVPR. (2018) 3339–3348
- Zhan, F., Xue, C., Lu, S.: Ga-dan: Geometry-aware domain adaptation network for scene text detection and recognition. In: ICCV. (2019) 9105–9115
- 34. Bhunia, A.K., Das, A., Bhunia, A.K., Kishore, P.S.R., Roy, P.P.: Handwriting recognition in low-resource scripts using adversarial learning. In: CVPR. (2019)
- 35. Lee, D.H.: Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In: Workshop, ICML. Volume 3. (2013)
- Choi, J., Jeong, M., Kim, T., Kim, C.: Pseudo-labeling curriculum for unsupervised domain adaptation. arXiv preprint arXiv:1908.00262 (2019)
- Chen, M., Weinberger, K.Q., Blitzer, J.: Co-training for domain adaptation. In: Advances in neural information processing systems. (2011) 2456–2464
- Saito, K., Ushiku, Y., Harada, T.: Asymmetric tri-training for unsupervised domain adaptation. arXiv preprint arXiv:1702.08400 (2017)
- Shu, R., Bui, H.H., Narui, H., Ermon, S.: A dirt-t approach to unsupervised domain adaptation. arXiv preprint arXiv:1802.08735 (2018)
- Zhang, W., Ouyang, W., Li, W., Xu, D.: Collaborative and adversarial network for unsupervised domain adaptation. In: CVPR. (2018) 3801–3809
- 41. Inoue, N., Furuta, R., Yamasaki, T., Aizawa, K.: Cross-domain weakly-supervised object detection through progressive domain adaptation. In: CVPR. (2018)
- Ózgen, A.C., Fasounaki, M., Ekenel, H.K.: Text detection in natural and computergenerated images. In: 2018 26th Signal Processing and Communications Applications Conference (SIU), IEEE (2018) 1–4
- Wang, W., Xie, E., Song, X., Zang, Y., Wang, W., Lu, T., Yu, G., Shen, C.: Efficient and accurate arbitrary-shaped text detection with pixel aggregation network. In: ICCV. (2019) 8440–8449
- Xing, L., Tian, Z., Huang, W., Scott, M.R.: Convolutional character networks. In: ICCV. (2019) 9126–9136