

# Stratified Domain Adaptation: A Progressive Self-Training Approach for Scene Text Recognition

Kha Nhat Le, Hoang-Tuan Nguyen, Hung Tien Tran, Thanh Duc Ngo\*  
University of Information Technology, VNU-HCM, Vietnam  
Vietnam National University, Ho Chi Minh City, Vietnam  
{20520208, 20520344, 19521587}@gm.uit.edu.vn, thanhnd@uit.edu.vn

## Abstract

*Unsupervised domain adaptation (UDA) has become increasingly prevalent in scene text recognition (STR), especially where training and testing data reside in different domains. The efficacy of existing UDA approaches tends to degrade when there is a large gap between the source and target domains. To deal with this problem, gradually shifting or progressively learning to shift from domain to domain is the key issue. In this paper, we introduce the Stratified Domain Adaptation (StrDA) approach, which examines the gradual escalation of the domain gap for the learning process. The objective is to partition the target data into subsets so that the progressively self-trained model can adapt to gradual changes. We stratify the target data by evaluating the proximity of each data sample to both the source and target domains. We propose a novel method for employing domain discriminators to estimate the out-of-distribution and domain discriminative levels of data samples. Extensive experiments on benchmark scene-text datasets show that our approach significantly improves the performance of baseline (source-trained) STR models. The source code is available at <https://github.com/KhaLee2307/StrDA>.*

## 1. Introduction

Although recent STR models have shown impressive performance, they are typically trained exclusively on labeled synthetic data. Baek *et al.* [4] have emphasized the significance of training STR models on real data, asserting its greater importance compared to synthetic data. However, collecting labeled real data poses a considerable challenge because of its high cost and time-intensive nature. Some efforts have been dedicated to generating synthetic data that closely resembles real data. The problem remains challenging due to the domain gap. Significant performance degradation occurs when a model trained with synthetic data is

applied to real data due to the substantial disparity between data distributions across domains. To address this problem, domain adaptation approaches are proposed to reduce distribution offsets. Especially, learning approaches that involve gradually shifting or progressively learning have demonstrated more notable improvement than learning directly from one source domain into another target domain.

In addition, while labeled real data is scarce, unlabeled real data is abundant and easily collectible. Several approaches used self-training methods to harness both labeled synthetic and unlabeled real data, with the aim of improving model performance [4, 23, 32, 33, 41, 52, 53]. This has demonstrated considerable effectiveness in improving the performance of the model and enabling the model to leverage the latent knowledge from unlabeled data points. However, it comes with several drawbacks due to its inherent instability. There is no explicit guarantee of the accuracy of the pseudo-labeling, which could cause the model to degrade. As the classification model weakens or the gap between the source and target domains increases, this phenomenon becomes more pronounced, making it challenging to control the upper bound on self-training errors [28].

In this work, we propose the Stratified Domain Adaptation (StrDA) approach by leveraging the gradual escalation of the domain gap to effectively address the discrepancy between the source and target domains. We partition the learning set from the target domain into smaller subsets, such that the domain gap of each subset compared to the source domain progressively increases. This way, the model can gradually adapt to domain changes and improve its performance. To evaluate the proximity of each data sample to the source and target domains, we propose the Harmonic Domain Gap Estimator (HDGE), which employs a pair of discriminators. Each discriminator evaluates the out-of-distribution (OOD) levels for each data point, with particular reference to the source or target domain. Then, these two OOD-level evaluations are passed through a harmonic function to estimate the distance from the data to the source domain. This means that data points that are situated

\*Corresponding author

near the intersection of the two domains exhibit a minimal distance, while data points that are outside the distribution of both domains display an exceedingly large distance. The harmonic evaluation function provides a more precise assessment of out-of-distribution levels, ensuring a more reliable assessment of OOD data points.

We summarize our contributions as follows:

- We introduce a progressive self-training domain adaptation approach for scene text recognition, which helps improve the model’s performance by utilizing unlabeled data with high-quality pseudo-labels. We propose the Harmonic Domain Gap Estimator method for stratifying domain gaps by analyzing the gradual escalation of domain gaps, which plays a crucial role in the effectiveness of domain adaptation.
- Extensive experiments are conducted on six benchmark datasets (IIIT [37], SVT [58], IC13 [26], IC15 [25], SVTP [43], CUTE [44]) and five additional datasets (COCO [57], Uber [67], ArT [11], ReCTS [66], Union14M [24]) to assess the performance of the proposed approach. It leads to a significant improvement in the performance of various existing STR models. This paves the way for recognizing text without incurring human annotation costs, particularly in cases where labeled real data is limited.

## 2. Related Work

### 2.1. Scene Text Recognition

In general, Scene Text Recognition (STR) is treated as a sequence prediction task that utilizes sequence modeling to leverage robust visual features for recognition. The CTC-based [18] decoder methods [2, 15, 20, 46] aim to maximize the probability of all possible paths for the final prediction, achieving a balance between accuracy and efficiency. Attention-based decoder methods [8, 10, 29, 31, 47, 48, 59, 60, 63, 65] utilize a visual query to localize the position of each character via an attention mechanism with the idea inspired from NLP community [9, 56]. This approach has demonstrated robustness in precision, albeit with high computational costs. Furthermore, some studies [13, 16, 38, 70] have shown the effectiveness of integrating the external language model to capture text semantics.

Most of the STR methods mentioned above have achieved remarkable results on common benchmarks, even when trained solely on synthetic datasets in a supervised manner. However, there is a significant domain gap between synthetic and real data. Jiang *et al.* [24] identified that STR is far from being solved by analyzing the numerous challenges associated with real data from a data-oriented perspective.

### 2.2. Domain Adaptation for Scene Text Recognition

Unsupervised Domain Adaptation (UDA) is a technique designed to enhance the performance of models trained on the source domain when applied to the target domain. One widely used approach in UDA is self-training [30]. The self-training process involves training a model with labeled data, then using this model to generate pseudo-labels for unlabeled data, which are subsequently used to retrain the target model. However, pseudo-labeling (PL) is often suboptimal due to erroneous predictions from poorly calibrated models, which can negatively impact training efficiency [1, 42]. Recent works [42, 45, 62] have focused on reducing PL errors in self-learning for general tasks and have demonstrated its effectiveness in more specific applications as well.

Recently, various studies have focused on refining the PL processes to harness labeled synthetic data and unlabeled real data in STR. Baek *et al.* [4] explored multiple ways to enhance STR models by using pseudo-labels. Patel *et al.* [40] introduced an uncertainty-based label selection strategy for STR by utilizing beam-search inference. Fang *et al.* [16] proposed the ensemble self-training strategy by treating the iterative predictions as an ensemble. Li *et al.* [32] introduced the Adaptive Distribution Regularizer to bridge the domain gap and achieved remarkable performance in cross-domain adaptation with both scene text and handwritten text. In another alternative perspective, several works [5, 34, 35, 55, 64, 68, 69] have recently suggested domain adaptation techniques to learn feature discrepancy between source and target domains. Zhang *et al.* [68] employed an Adversarial Sequence-to-Sequence Domain Adaptation (ASSDA) network, which could adaptively align the coarse global-level and fine-grained character-level representation across domains in an adversarial manner. Liu *et al.* [34] introduced ProtoUDA, which enhances text recognition across various domains by using pseudo-labeled character features and parallel, complementary modules for class-level and instance-level alignment.

The inherent domain gap between labeled and unlabeled data mainly results in low-quality derived pseudo-labels. Although these methods offer promising results, they primarily focus on addressing the domain gap through direct adaptation. We take into account that the domain gap has a progressive tendency. Instead of directly adapting from the source to the target domain, we propose to leverage the gradual shift between domains in scene text recognition. Recent studies in other computer vision tasks have also presented approaches such as Gradual Adaptation [7, 14, 17, 21] and Easy-to-Hard Transfer [6, 61] for adapting to different domains. These approaches demonstrate significant performance enhancement compared to direct UDA.

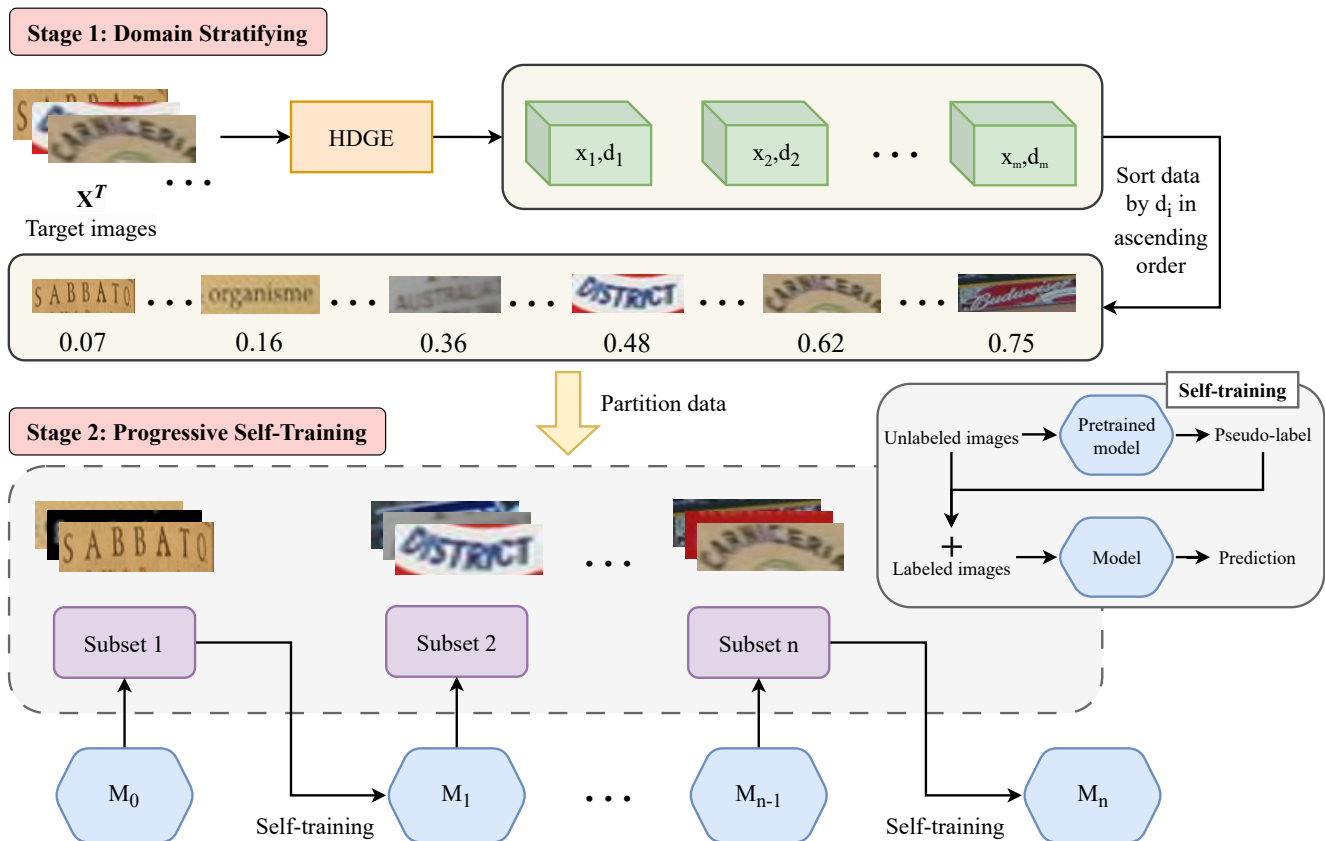


Figure 1. **The overall framework of our proposed Stratified Domain Adaptation (StrDA) for scene text recognition.** Our approach leverages labeled synthetic data and unlabeled real data, without human annotation. The entire process is divided into 2 stages: Domain Stratifying (partitioning the unlabeled real data into subsets satisfying Eq. (1)) and Progressive Self-Training.  $m$  represents the number of unlabeled data, and  $n$  serves as the hyper-parameter in Eq. (1).

### 3. Stratified Domain Adaptation

#### 3.1. Overview

In this work, our focus is on addressing the problem using two predefined datasets: one comprising labeled data samples from the source domain,  $S = \{(\mathbf{x}_i^S, \mathbf{y}_i^S)\}_{i=1}^{|S|}$ ; and, the other comprising unlabeled data samples from the target domain,  $T = \{\mathbf{x}_i^T\}_{i=1}^{|T|}$ . The goal of domain adaptation is to enhance the performance of the source-trained model by leveraging both  $S$  and  $T$ .

**Unsupervised Domain Adaptation (UDA).** To investigate the Stratified Domain Gap approach, we relied on the traditional UDA approach using *vanilla self-training* (denoted as ST). ST takes a source-trained model (referred to as the *baseline model*) to generate a pseudo-label for  $\mathbf{x}_i^T$ . Subsequently, the model is trained using the pseudo-labeled data combined with the labeled data from the source domain. Applying domain adaptation directly (using the entire dataset for a single self-training process) may encounter several disadvantages (Sec. 1). Instead, our approach em-

plloys a series of ST rounds with a sequence of target sub-domain data.

We first partition the unlabeled data into a sequence of equally-sized subsets  $T_1, T_2, T_3, \dots, T_n$ , where  $T_m = \{\mathbf{x}_i^T\}_{i=1}^{|T_m|}$ . By this, we assume that the domain gap between  $T_m$  and  $S$  is less than that between  $T_{m+1}$  and  $S$ :

$$\rho(S, T_m) \leq \rho(S, T_{m+1}), \quad \forall m \in (1, n) \quad (1)$$

where  $\rho(P, Q)$ <sup>1</sup> is a distance function between distributions  $P$  and  $Q$ .

To partition the data with respect to Eq. (1), we propose the Harmonic Domain Gap Estimator (HDGE) method to estimate the proximity of a data point  $\mathbf{x}_i \in T$  and the source domain  $S$ . Afterward, we arrange and partition the data that satisfy Eq. (1). We refer to the entire process as Stage 1- Domain Stratifying. After obtaining the subsets from stage 1, we sequentially apply ST to each subset. This process

<sup>1</sup> $p$  can be Kullback-Leibler Divergence (KL Divergence) or Wasserstein Distance

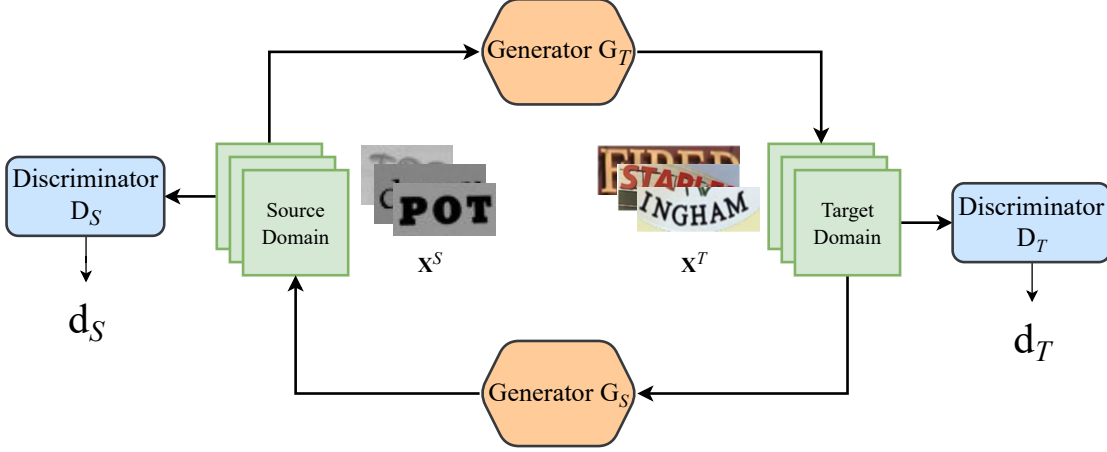


Figure 2. Our architecture consists of two mapping functions,  $G_T : S \rightarrow T$  and  $G_S : T \rightarrow S$ , along with associated adversarial discriminators,  $D_S$  and  $D_T$ . While  $G_S$  and  $G_T$  are tasked with translating images from one domain to another,  $D_S$  estimates the difference between an image and the data distribution of the source domain  $S$ , and similarly,  $D_T$  does so for the target domain  $T$ .

is referred to as Stage 2-Progressive Self-Training. The entire Stratified Domain Adaptation approach consists of two stages, as described in Fig. 1.

### 3.2. Stage 1: Domain Stratifying

Given  $S$  and  $T$ , we introduce HDGE to assess the proximity of a data point  $\mathbf{x}_i^T \in T$  and  $S$ , denoted as  $d_i$ . Lower  $d_i$  indicates that  $\mathbf{x}_i^T$  is closer to the source. After assigning  $d_i$  to each data point  $\mathbf{x}_i^T$ , we arrange the data in ascending order of  $d_i$  and then partition them into  $n$  subsets with equal size,  $T_m = \left\{ \mathbf{x}_i^T \right\}_{i=1}^{|T_m|}$ , for progressive self-training.

**Harmonic Domain Gap Estimator (HDGE)** uses a pair of discriminators, one for the source domain and the other for the target domain ( $D_S$  and  $D_T$ ). Each discriminator evaluates the out-of-distribution (OOD) levels for each data point. By synthesizing the outputs of the two discriminators, we can determine whether the data is in-domain (near source or target) or out of both distributions. We denote these two OOD levels as  $d_S$  and  $d_T$ . To calculate the  $d_i$  for  $\mathbf{x}_i^T$ , we use the formula:

$$d_i = \frac{(1 + \beta^2) \cdot d_S(\mathbf{x}_i^T) \cdot d_T(\mathbf{x}_i^T)}{\beta^2 \cdot d_S(\mathbf{x}_i^T) + d_T(\mathbf{x}_i^T)} \quad (2)$$

where  $0 \leq \beta < 1$ , we tend to bias the data towards smaller  $d_S(\mathbf{x}_i^T)$ , meaning closer to the source domain. This aligns with the condition Eq. (1).

With the designed  $d_i$ -computation function as above, we aim to arrange the data for the progressive self-training process with the following prioritization:

1.  $\mathbf{x}_i$  situated at intersection of two distributions ( $d_S$  and  $d_T$  are small)
2.  $\mathbf{x}_i$  closer to the source domain (small  $d_S$ , large  $d_T$ )

3.  $\mathbf{x}_i$  closer to the target domain (small  $d_T$ , large  $d_S$ )
4.  $\mathbf{x}_i$  that is out of two distributions ( $d_S$  and  $d_T$  are large)

To create a pair of discriminators  $D_S$  and  $D_T$  with the ability to assess out-of-distribution (OOD) levels effectively, we designed a learning strategy inspired by [72]. As illustrated in Fig. 2, in addition to the two discriminators  $D_S$  and  $D_T$ , we also utilize two generators:  $G_T$  translates images from the source domain to the target domain ( $G_T : S \rightarrow T$ ), and  $G_S$  performs a similar task from the target domain to the source domain ( $G_S : T \rightarrow S$ ).

While *generators* strive to learn how to represent from one domain to another, *discriminators* learn to distinguish between images generated by the generator and *real* images. Through adversarial learning,  $G_S$  and  $G_T$  will improve image generation, consequently enhancing the discriminative abilities of  $D_S$  and  $D_T$ . As a result, when a new data point  $\mathbf{x}_i$  is introduced, the discriminator pair accurately assesses out-of-distribution levels ( $d_S$  and  $d_T$ ).

Given training samples  $\{\mathbf{x}_i^S\}_{i=1}^{|S|}$  where  $\mathbf{x}_i^S \in S$  and  $\{\mathbf{x}_i^T\}_{i=1}^{|T|}$  where  $\mathbf{x}_i^T \in T$ , the data distribution is indicated as  $\mathbf{x}^S \sim p_{\text{data}}(\mathbf{x}^S)$  and  $\mathbf{x}^T \sim p_{\text{data}}(\mathbf{x}^T)$ . The adversarial loss for the mapping function  $G_T : S \rightarrow T$  and its discriminator  $D_T$  is expressed as follows:

$$L_{GAN}(G_T, D_T, S, T) = \mathbb{E}_{\mathbf{x}^T \sim p_{\text{data}}(\mathbf{x}^T)} [\log D_T(\mathbf{x}^T)] + \mathbb{E}_{\mathbf{x}^S \sim p_{\text{data}}(\mathbf{x}^S)} [\log(1 - D_T(G_T(\mathbf{x}^S)))] \quad (3)$$

where  $G_T$  attempts to generate images  $G_T(\mathbf{x}^S)$  that resemble images from domain  $T$ , while the objective of  $D_T$  is to differentiate between translated samples  $G_T(\mathbf{x}^S)$  and real samples  $\mathbf{x}^T$ .  $G_T$  strives to minimize this objective against the adversary  $D_T$  that seeks to maximize it, i.e.  $\min_{G_T} \max_{D_T} L_{GAN}(G_T, D_T, S, T)$ . We

use a similar adversarial loss for the mapping function  $G_S : T \rightarrow S$  and its discriminator  $D_S$  as well: *i.e.*  $\min_{G_S} \max_{D_S} L_{GAN}(G_S, D_S, T, S)$

After training, we obtain a pair of discriminators,  $D_S$  and  $D_T$  with the ability to estimate the domain gap  $d_i$  for data  $x_i^T$  using Eq. (2).

### 3.3. Stage 2: Progressive Self-Training

At the end of Stage 1, we have  $n$  subsets for Stage 2-Progressive Self-Training. As demonstrated in Fig. 1, we will conduct self-training (ST) sequentially on each set of sub-domain data  $T_i$ . The entire learning process is described in Algorithm 1.

---

#### Algorithm 1 Progressive Self-Training ST

---

**Require:** Labeled images  $(X, Y) \in S$  and sequence of unlabeled image subsets  $T_1, T_2, T_3, \dots, T_n (T_i \in T)$

- 1: Train STR model  $M(\cdot, \theta_0)$  with  $(X, Y)$  using Eq. (4).
  - 2: **for** iteration  $i = 1, 2, \dots, n$  **do**
  - 3:  $T_i \rightarrow M(\cdot, \theta_{i-1}) \rightarrow V_i$  (pseudo-labels) and  $m_i$  (average confidence-scores)
  - 4: Update  $\theta_i$  with  $(X, Y), (T_i, V_i), m_i$  using Eq. (5)
  - 5: **end for**
- 

Given the input image  $x^L$  and the character sequence of the ground truth  $y^L = y_1^L, \dots, y_k^L$ , the STR model  $M(\cdot; \theta)$  outputs a vector sequence  $\mathbf{p}^L = M(x^L; \theta) = p_1^L, \dots, p_k^L$ . Cross-entropy loss is employed to train the STR model:

$$L_r(x^L, y^L) = \frac{1}{k} \sum_{i=1}^k \log p_i^L(y_i^L | x^L) \quad (4)$$

where  $p_i^L(y_i^L)$  represents the predicted probability of the output being  $y_i^L$  at time step  $t$  and  $k$  is the sequence length.

In each ST round, after obtaining labeled data and pseudo-labeled data, we proceed to train the STR model  $M(\cdot; \theta)$  to minimize the objective function:

$$L(\phi) = \frac{1 - m_i}{|S|} \sum_{x^S \in S} L_r(x^S; y^S) + \frac{m_i}{|T_i|} \sum_{x^{T_i} \in T_i} L_r(x^{T_i}; y^{T_i}) \quad (5)$$

where  $m_i$  is the mean (average) of confidence scores when generating pseudo-labels for the unlabeled image subset  $T_i$ .  $m_i$  serves as an *adaptive controller*.

### 3.4. Additional Training Techniques

**Label Sharpening.** We "sharpen" the soft labels to encourage the model to update its parameters. Consequently, during the training process in Stage 2, we utilize the model's predictions on unlabeled data as definitive pseudo-labels rather than relying on their probabilities.

**Regularization.** Regularization is a significant factor in self-training. Without regularization, the model is not incentivized to change during self-training [28]. Therefore, we also incorporate it into our model training process.

**Data Augmentation.** We apply multiple augmentation strategies on both geometry transformations and color jitter, which are borrowed from RandAugment [12].

## 4. Experiments

### 4.1. Datasets

Our work focuses on addressing the domain gap problem between the source domain, which is **synthetic** data, and the target domain, which is **real** data in *scene text recognition*.

Experiments are conducted according to the setup of [3] to ensure a fair comparison. We used two types of data during the training process: synthetic data (with SynthText (ST) [19] and MJSynth (MJ) [22]) and real data without labels. Concretely, we collect public real datasets, including ArT [11], COCO-Text (COCO) [57], LSVT [54], MLT19 [39], OpenVINO [27], RCTW17 [49], ReCTS [66], Uber-Text (Uber) [67], TextOCR [50], and **discard their labels** to formulate the set of real data without labels. In addition, we exclude vertical text (height > width) and images whose width is greater than 25 times the height. As a result, we have 16 million labeled synthetic data and 2 million unlabeled real data for training, denoted as **real unlabeled data (2M RU)**.

For evaluation, six standard benchmark datasets, including IIIT 5k-word (IIIT) [37], Street View Text (SVT) [58], ICDAR 2013 (IC13) [26], ICDAR 2015 (IC15) [25], SVT-Perspective (SVTP) [43], and CUTE80 (CUTE) [44], are used. Note that IC13 and IC15 have two versions of their respective test splits commonly used in the literature: 857 and 1,015 for IC13; 1,811 and 2,077 for IC15.

In order to achieve a comprehensive comparison, we expand our evaluation to encompass five larger and more challenging datasets: COCO-Text (COCO) [57], Uber-Text (Uber) [67], ArT [11], ReCTS [66], and Union14M [24] (Artistic, Contextless, Curve, General).

### 4.2. Evaluation Metrics

Following standard conventions [3], we present word-level accuracy for each dataset. Furthermore, to provide a thorough evaluation of the models concerning their recognition performance on both *regular* and *irregular* text, as per [4], we introduce an average score denoted "Avg." This score represents the accuracy across the combined set of samples from all six benchmark datasets (IIIT<sub>3000</sub>, SVT<sub>647</sub>, IC13<sub>1015</sub>, IC15<sub>2077</sub>, SVTP<sub>645</sub>, and CUTE<sub>288</sub>).

Table 1. **Word accuracy on six scene-text benchmarks and five additional datasets.** The number of words in each dataset is listed with its name. "(baseline)" means models trained only on synthetic data. We present the results of domain adaptation approaches for the baseline model. "ST" refers to traditional unsupervised domain adaptation (vanilla self-training). "StrDA<sub>HDGE</sub>" is Stratified Domain Adaptation methods using Harmonic Domain Gap Estimator. The numbers denoted by "Δ" in **green** indicate improvements over each dataset. With each dataset (*i.e.*, each column), the best result is shown in **bold**. With each baseline model on a dataset, the best result is shown with an underline. The performance of baseline models is substantially enhanced by the proposed methods.

Type	Method	Common Benchmarks							Additional Datasets				
		IIT	SVT	IC13	IC15	SVTP	CUTE	Avg.	COCO	Uber	ArT	ReCTS	Union14M
		3,000	647	857	1,811	645	288		9,825	80,418	35,149	2,592	403,379
CTC	CRNN (baseline)	92.5	86.4	92.0	71.3	75.2	83.3	84.4	49.5	34.6	59.5	77.1	43.3
	+ ST	<u>93.7</u>	87.6	92.2	72.9	75.5	<u>84.7</u>	85.5	51.4	35.9	60.7	79.8	46.2
	Δ	<b>+1.2</b>	<b>+1.2</b>	<b>+0.2</b>	<b>+1.6</b>	<b>+0.3</b>	<b>+1.4</b>	<b>+1.1</b>	<b>+1.9</b>	<b>+1.3</b>	<b>+1.2</b>	<b>+2.7</b>	<b>+2.9</b>
	+ StrDA <sub>HDGE</sub>	93.4	<u>89.0</u>	<u>93.1</u>	<u>74.0</u>	<u>77.1</u>	84.4	<u>86.0</u>	<u>53.0</u>	<u>36.8</u>	<u>60.9</u>	<u>81.0</u>	<u>47.8</u>
	Δ	<b>+0.9</b>	<b>+2.6</b>	<b>+1.1</b>	<b>+2.7</b>	<b>+1.9</b>	<b>+1.1</b>	<b>+1.6</b>	<b>+3.5</b>	<b>+2.2</b>	<b>+1.4</b>	<b>+3.9</b>	<b>+4.5</b>
Attention	TRBA (baseline)	96.2	93.7	95.8	81.9	86.1	91.0	91.0	62.5	39.0	69.0	82.8	56.6
	+ ST	97.1	94.0	96.1	82.5	90.1	92.4	92.0	65.5	40.9	70.9	84.8	60.4
	Δ	<b>+0.9</b>	<b>+0.3</b>	<b>+0.3</b>	<b>+0.6</b>	<b>+4.0</b>	<b>+1.4</b>	<b>+1.0</b>	<b>+3.0</b>	<b>+1.9</b>	<b>+1.9</b>	<b>+2.0</b>	<b>+3.8</b>
	+ StrDA <sub>HDGE</sub>	<u>97.2</u>	<u>95.2</u>	<b>96.5</b>	<b>84.5</b>	<u>90.7</u>	<b>94.4</b>	<u>92.8</u>	<u>68.6</u>	<u>42.7</u>	<u>72.2</u>	<b>85.8</b>	<b>64.2</b>
	Δ	<b>+1.0</b>	<b>+1.5</b>	<b>+0.7</b>	<b>+2.6</b>	<b>+4.6</b>	<b>+3.4</b>	<b>+1.8</b>	<b>+6.1</b>	<b>+3.7</b>	<b>+3.2</b>	<b>+3.0</b>	<b>+7.6</b>
LM	ABINet (baseline)	97.0	95.2	95.6	82.3	89.5	90.3	91.8	63.2	39.5	68.9	82.6	55.7
	+ ST	97.4	96.3	<u>96.4</u>	83.9	<b>91.0</b>	92.0	92.8	68.7	42.3	71.2	84.7	61.4
	Δ	<b>+0.4</b>	<b>+1.1</b>	<b>+0.8</b>	<b>+1.6</b>	<b>+1.5</b>	<b>+1.7</b>	<b>+1.0</b>	<b>+5.5</b>	<b>+2.8</b>	<b>+2.3</b>	<b>+2.1</b>	<b>+5.7</b>
	+ StrDA <sub>HDGE</sub>	<b>97.8</b>	<b>96.9</b>	96.0	<u>84.4</u>	<b>91.0</b>	<b>94.4</b>	<b>93.2</b>	<b>69.7</b>	<b>44.2</b>	<u>71.6</u>	<u>85.0</u>	<u>62.9</u>
	Δ	<b>+0.8</b>	<b>+1.7</b>	<b>+0.4</b>	<b>+2.1</b>	<b>+1.5</b>	<b>+4.1</b>	<b>+1.4</b>	<b>+6.5</b>	<b>+4.7</b>	<b>+2.7</b>	<b>+2.4</b>	<b>+7.2</b>

### 4.3. Implementation Details

Three STR models, CRNN [46], TRBA [3] and ABINet [16], are employed to assess the effectiveness of the proposed framework using their default configurations. We trained the *baseline* STR models in a fully supervised manner using the synthetic dataset (MJ+ST). Our reproduced results from supervised training exceed those reported in the original papers [3, 16, 46]. Besides the adopted augmentation techniques [12], we trained the STR models for more iterations (300K).

For Stage 1 (Domain Stratifying), the Harmonic Domain Gap Estimator (HDGE) utilizes a generative adversarial network. Details are in the supplementary materials.

For Stage 2 (Progressive Self-Training), we adopt the AdamW [36] optimizer (weight\_decay 0.005). We also use the one-cycle learning rate scheduler [51] with a maximum learning rate of 0.0005. Our training batch size is fixed at 128, and the total number of iterations is 50K.

To demonstrate the effectiveness of our StrDA approach compared to traditional unsupervised domain adaptation (vanilla self-training ST) (Sec. 3.1), we conducted all experiments with the same protocols. All experiments were performed on an NVIDIA RTX A5000 (24GB VRAM).

### 4.4. Results and Analysis

As illustrated in Tab. 1, both domain adaptation methods (ST and StrDA<sub>HDGE</sub>) surpass the baseline models across eleven public benchmarks. Despite relying solely on additional unlabeled real data and self-training with pseudo-labels, the experiments remarkably enhanced the STR model's performance on both *regular* and *irregular* datasets. These remarkable results emphasize the importance of integrating real images into training STR models.

Notably, the StrDA<sub>HDGE</sub> method applied to all three baseline models of STR outperforms vanilla self-training (ST). We observed that ST does not perform well without domain sequences, although it shows a slight improvement over the source-trained model (improved by 1.1% for CRNN and 1% for both TRBA and ABINet on Avg.). Our progressive self-training framework shows strong effectiveness by partitioning and organizing data according to the progressive increase in domain gap.

Specifically, CRNN, TRBA, and ABINet exhibit remarkable improvements on Avg. (**+1.6%**, **+1.8%**, and **+1.4%**) when applying StrDA<sub>HDGE</sub>. Furthermore, with large and challenging datasets such as Union14M, the StrDA<sub>HDGE</sub> method demonstrates exceptional effectiveness (**+4.5%**, **+7.6%**, and **+7.2%**). These results demonstrate the gen-

Table 2. Comparison with other domain adaptation methods in the STR task. Our method significantly enhances the performance of the STR models, surpassing other existing approaches. Additionally, it can be integrated with other methods to achieve even greater efficacy.

	Method	Labeled Dataset	Unlabeled Dataset	Regular Text				Irregular Text			
				IIIT	SVT	IC13		IC15	SVTP	CUTE	
				3000	647	857	1015	1811	2077	645	288
Published Results	TRBA-FEDS [41]	MJ+ST	Amazon_book_cover	92.2	92.1	96.5	95.3	83.8	80.9	84.0	79.0
	TRBA-Seq-UPS [40]	MJ+ST	276K RU	92.7	88.6	-	92.2	-	76.9	78.0	84.4
	TRBA-cr [71]	MJ+ST	10.6M RU	96.5	96.3	<b>98.3</b>	-	<b>89.3</b>	-	<b>93.3</b>	93.4
	ABINet-st [16]	MJ+ST	Uber-Text	96.8	94.9	97.3	-	87.4	-	90.1	93.4
	ABINet-est [16]	MJ+ST	Uber-Text	97.2	95.5	97.7	-	86.9	-	89.9	94.1
Our Results	TRBA-cr (reproduce)	MJ+ST	2M RU	97.3	95.1	97.2	96.2	88.1	84.0	90.5	93.8
	<b>TRBA-StrDA<sub>HDGE</sub></b>	MJ+ST	2M RU	97.2	95.2	97.4	96.5	88.4	<b>84.5</b>	90.7	<b>94.4</b>
	<b>TRBA-StrDA<sub>HDGE</sub> w/ cr</b>	MJ+ST	2M RU	97.3	96.1	97.6	<b>96.7</b>	88.7	<b>84.5</b>	90.9	<b>94.4</b>
	<b>ABINet-StrDA<sub>HDGE</sub></b>	MJ+ST	2M RU	<b>97.8</b>	<b>96.9</b>	97.0	96.0	88.6	84.4	91.0	<b>94.4</b>

Table 3. Effect of our proposed HDGE. Compared to using DD, StrDA with HDGE yields better results.

Method	IIIT	SVT	IC13	IC15	SVTP	CUTE	Avg.
	3000	647	1015	2077	645	288	
CRNN-StrDA <sub>DD</sub>	93.3	88.9	92.5	73.4	76.7	83.0	85.7
CRNN-StrDA <sub>HDGE</sub>	<u>93.4</u>	<u>89.0</u>	<u>93.1</u>	<u>74.0</u>	<u>77.1</u>	<u>84.4</u>	<u>86.0</u>
TRBA-StrDA <sub>DD</sub>	<u>97.6</u>	94.7	96.1	83.6	89.9	93.1	92.5
TRBA-StrDA <sub>HDGE</sub>	97.2	<u>95.2</u>	<b>96.5</b>	<b>84.5</b>	<u>90.7</u>	<b>94.4</b>	<u>92.8</u>
ABINet-StrDA <sub>DD</sub>	<b>97.8</b>	96.8	<u>96.2</u>	84.1	<b>91.0</b>	93.4	93.0
ABINet-StrDA <sub>HDGE</sub>	<b>97.8</b>	<b>96.9</b>	96.0	<u>84.4</u>	<b>91.0</b>	<b>94.4</b>	<b>93.2</b>

eralizability of our proposed methods, as StrDA<sub>HDGE</sub> is effective across various STR models, including CTC-based, Attention-based, and LM-based models.

#### 4.5. Comparison with Other Methods

In Tab. 2, we perform a comparative analysis of our proposed method with other unsupervised domain adaptation methods for scene text recognition. As we reimplemented the TRBA-cr method, the reproduced results were slightly different from those reported in the original paper [71]. This is because we used less data (2M compared to 10.6M) while keeping all other settings the same.

TRBA-StrDA<sub>HDGE</sub> achieved superior results in most datasets. Furthermore, when combining both methods, TRBA-StrDA<sub>HDGE</sub> with *cr*, the performance improved beyond what was achieved independently by either method. For other methods compared on the same backbone, StrDA<sub>HDGE</sub> consistently demonstrated superior performance. The previous works are commendable. It is noteworthy that our framework can be conceptualized as a series of domain adaptation rounds. Consequently, integrating advanced techniques into stage 2 of our framework would be highly advantageous.

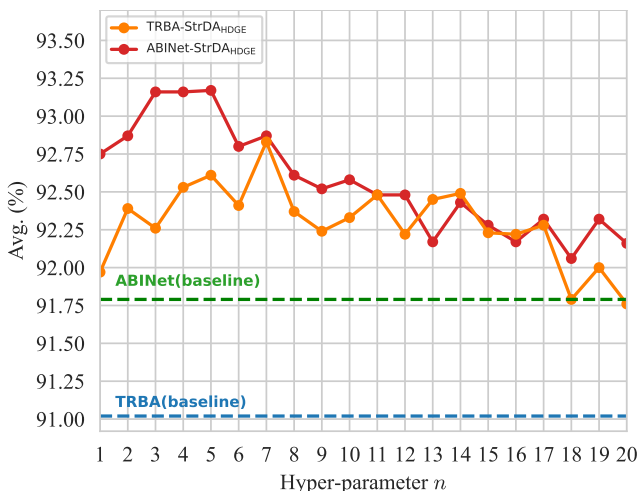


Figure 3. Ablation study on Hyper-parameter  $n$  in Eq. (1).

#### 4.6. Ablation Study

In this section, we conduct a series of ablation experiments. TRBA and ABINet are used in all experiments due to their superior performance. The dataset used consists of 2M RU. Additional experiments are provided in the Supp.

##### 4.6.1 Ablation Study on Hyper-parameter $n$ in Eq. (1)

In this section, we proceed to compare the performance of StrDA<sub>HDGE</sub> with different numbers of subsets (with the hyper-parameter  $\beta = 0.9$ , and every subset has the same size). The case where  $n = 1$  corresponds to vanilla self-training (ST) as described in Sec. 3.1.

According to Fig. 3, StrDA<sub>HDGE</sub> proves to be more effective for both methods as the hyper-parameter  $n$  increases, reaching **92.83%** at  $n = 7$  for TRBA and **93.17%** at  $n = 5$

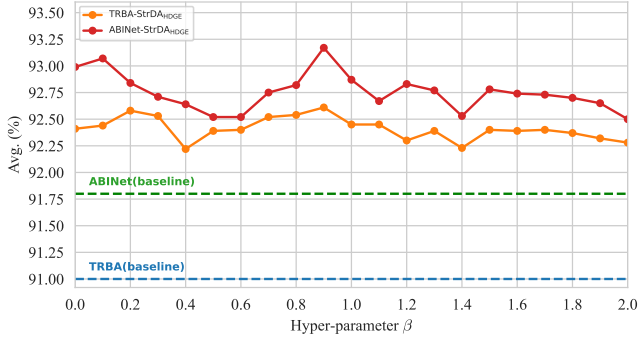


Figure 4. Ablation study on Hyper-parameter  $\beta$  in Eq. (2)

for ABINet. This phenomenon is understandable due to the reduction in domain disparity resulting from smaller data partitions, thereby facilitating the adaptability of the model. However, with larger values  $n$ , the performance of the method tends to saturate and eventually decline, suggesting the need for a judicious choice of  $n$ . Future work should seek generalized approaches to determine  $n$  and the number of data samples in each subset.

#### 4.6.2 Ablation Study on Hyper-parameter $\beta$ in Eq. (2)

We conducted experiments (with the hyper-parameter  $n = 5$ ) to observe the influence of the hyper-parameter  $\beta$  in Eq. (2) on the effectiveness of the  $\text{StrDA}_{\text{HDGE}}$  method. We are adjusting the value of  $\beta$  either higher or lower, causing HDGE to exhibit less or more bias towards the source domain. When  $\beta = 0$ ,  $d_i = d_S$ ,  $\text{StrDA}_{\text{HDGE}}$  only uses information from the source domain. In this case, as shown in Fig. 4,  $\text{StrDA}$  demonstrates fairly good effectiveness (92.41% for TRBA and 92.99% for ABINet).

However, incorporating information from both the source and target directions leads to significantly higher performance (92.61% and 93.17% for  $\beta = 0.9$ ). This supports our suggestion in Eq. (2). It also reinforces our claim that stratifying domain gaps using information from both source and target domains contributes to overall effectiveness.

#### 4.6.3 Alternative Domain Gap Estimators

We conducted additional experiments with an alternative gap estimator, called Domain Classifier (DD), to assess the domain gap  $d_i$  (Sec. 3.2). DD employs a binary classifier  $f(\mathbf{x}_i; \phi)$  with a feature extractor from the baseline model combined with a fully connected layer at the final layer. DD is trained using raw images from  $S$  (assigned as class 0) and  $T$  (assigned as class 1). Next, we assign  $d_i = f(\mathbf{x}_i; \phi)$ . By learning distinctive features from the two domains, the discriminator identifies whether a data point is closer to the source or target domain, corresponding to a smaller or larger distance from the source.



Figure 5. The  $\text{StrDA}$  partitions the data from the target domain into three distinct subsets, with the disparity across domains gradually rising, as shown in the image. The next two lines depict the pseudo-labels employed in the self-training process of ST and  $\text{StrDA}_{\text{HDGE}}$ , respectively. The pseudo-labels generated by ST are prone to noise as the extent of the domain gap escalates. On the other hand,  $\text{StrDA}_{\text{HDGE}}$  produces pseudo-labels with higher accuracy. The STR model used for the example is TRBA.

As shown in Tab. 3,  $\text{StrDA}_{\text{HDGE}}$  yields better results compared to  $\text{StrDA}_{\text{DD}}$ . We note that DD treats data points situated in the intersection and those outside both distributions similarly, with the same  $d_i$ . This leads to poor discrimination in the self-learning process.

## 5. Conclusion

In this paper, we propose the Stratified Domain Adaptation ( $\text{StrDA}$ ) approach, a progressive self-training framework for scene text recognition. By leveraging the gradual escalation of the domain gap with the Harmonic Domain Gap Estimator (HDGE), we propose partitioning the target domain into a sequence of ordered subsets to progressively reduce the domain gap between each and the source domain. Progressive self-training is then applied sequentially to these subsets. Extensive experiments on STR benchmarks demonstrate that our approach enables the baseline STR models to progressively adapt to the target domain. This approach significantly improves the performance of the baseline model without using any human-annotated data and shows its superior effectiveness compared to existing UDA methods for the scene text recognition task.

## Acknowledgements

This research is funded by University of Information Technology - VNUHCM under grant number D4-2024-03.



## References

- [1] Eric Arazo, Diego Ortego, Paul Albert, Noel E O'Connor, and Kevin McGuinness. Pseudo-labeling and confirmation bias in deep semi-supervised learning. In *2020 International joint conference on neural networks (IJCNN)*, pages 1–8. IEEE, 2020. [2](#)
- [2] Rowel Atienza. Vision transformer for fast and efficient scene text recognition. In *International conference on document analysis and recognition*, pages 319–334. Springer, 2021. [2](#)
- [3] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwal-suk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4715–4723, 2019. [5](#), [6](#)
- [4] Jeonghun Baek, Yusuke Matsui, and Kiyoharu Aizawa. What if we only use real datasets for scene text recognition? toward scene text recognition with fewer labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3113–3122, 2021. [1](#), [2](#), [5](#)
- [5] Yen-Cheng Chang, Yi-Chang Chen, Yu-Chuan Chang, and Yi-Ren Yeh. Smile: Sequence-to-sequence domain adaptation with minimizing latent entropy for text image recognition. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 431–435. IEEE, 2022. [2](#)
- [6] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 627–636, 2019. [2](#)
- [7] Hong-You Chen and Wei-Lun Chao. Gradual domain adaptation without indexed intermediate domains. *Advances in neural information processing systems*, 34:8201–8214, 2021. [2](#)
- [8] Changxu Cheng, Peng Wang, Cheng Da, Qi Zheng, and Cong Yao. Lister: Neighbor decoding for length-insensitive scene text recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19541–19551, 2023. [2](#)
- [9] Jianpeng Cheng, Li Dong, and Mirella Lapata. Long short-term memory-networks for machine reading. *arXiv preprint arXiv:1601.06733*, 2016. [2](#)
- [10] Zhazhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. Focusing attention: Towards accurate text recognition in natural images. In *Proceedings of the IEEE international conference on computer vision*, pages 5076–5084, 2017. [2](#)
- [11] Chee Kheng Chng, Yuliang Liu, Yipeng Sun, Chun Chet Ng, Canjie Luo, Zihan Ni, ChuanMing Fang, Shuaitao Zhang, Junyu Han, Errui Ding, et al. Icdar2019 robust reading challenge on arbitrary-shaped text-rrc-art. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1571–1576. IEEE, 2019. [2](#), [5](#)
- [12] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 702–703, 2020. [5](#), [6](#)
- [13] Cheng Da, Peng Wang, and Cong Yao. Levenshtein ocr. In *European Conference on Computer Vision*, pages 322–338. Springer, 2022. [2](#)
- [14] Dengxin Dai and Luc Van Gool. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3819–3824. IEEE, 2018. [2](#)
- [15] 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. *arXiv preprint arXiv:2205.00159*, 2022. [2](#)
- [16] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7098–7107, 2021. [2](#), [6](#), [7](#)
- [17] Michael Gadermayr, Dennis Eschweiler, Barbara Mara Klinkhammer, Peter Boor, and Dorit Merhof. Gradual domain adaptation for segmenting whole slide images showing pathological variability. In *Image and Signal Processing: 8th International Conference, ICISP 2018, Cherbourg, France, July 2-4, 2018, Proceedings 8*, pages 461–469. Springer, 2018. [2](#)
- [18] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376, 2006. [2](#)
- [19] 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. [5](#)
- [20] Pan He, Weilin Huang, Yu Qiao, Chen Loy, and Xiaoou Tang. Reading scene text in deep convolutional sequences. In *Proceedings of the AAAI conference on artificial intelligence*, volume 30, 2016. [2](#)
- [21] DuoJun Huang, Jichang Li, Weikai Chen, Junshi Huang, Zhenhua Chai, and Guanbin Li. Divide and adapt: Active domain adaptation via customized learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7651–7660, 2023. [2](#)
- [22] 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. [5](#)
- [23] Klara Janouskova, Jiri Matas, Lluís Gomez, and Dimosthenis Karatzas. Text recognition-real world data and where to find them. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 4489–4496. IEEE, 2021. [1](#)
- [24] Qing Jiang, Jiapeng Wang, Dezhi Peng, Chongyu Liu, and Lianwen Jin. Revisiting scene text recognition: A data perspective. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 20543–20554, 2023. [2](#), [5](#)

- [25] 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. [2](#), [5](#)
- [26] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. In *2013 12th international conference on document analysis and recognition*, pages 1484–1493. IEEE, 2013. [2](#), [5](#)
- [27] Ilya Krylov, Sergei Nosov, and Vladislav Sovrasov. Open images v5 text annotation and yet another mask text spotter. In *Asian Conference on Machine Learning*, pages 379–389. PMLR, 2021. [5](#)
- [28] Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*, pages 5468–5479. PMLR, 2020. [1](#), [5](#)
- [29] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2231–2239, 2016. [2](#)
- [30] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, page 896. Atlanta, 2013. [2](#)
- [31] Junyeop Lee, Sungrae Park, Jeonghun Baek, Seong Joon Oh, Seonghyeon Kim, and Hwalsuk Lee. On recognizing texts of arbitrary shapes with 2d self-attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 546–547, 2020. [2](#)
- [32] Xiaoyu Li, Xiaoxue Chen, Zuming Huang, Lele Xie, Jingdong Chen, and Ming Yang. Fine-grained pseudo labels for scene text recognition. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5786–5795, 2023. [1](#), [2](#)
- [33] Zheng Li, Joshua Smith, and Sujoy Chakraborty. Domain adaptation in sequence-to-sequence scene text recognition. 2021. [1](#)
- [34] Xiao-Qian Liu, Xue-Ying Ding, Xin Luo, and Xin-Shun Xu. Protouda: Prototype-based unsupervised adaptation for cross-domain text recognition. *IEEE Transactions on Knowledge and Data Engineering*, 2023. [2](#)
- [35] Xiao-Qian Liu, Xue-Ying Ding, Xin Luo, and Xin-Shun Xu. Unsupervised domain adaptation via class aggregation for text recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, 33(10):5617–5630, 2023. [2](#)
- [36] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. [6](#)
- [37] Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. In *BMVC-British machine vision conference*. BMVA, 2012. [2](#), [5](#)
- [38] Byeonghu Na, Yoonsik Kim, and Sungrae Park. Multi-modal text recognition networks: Interactive enhancements between visual and semantic features. In *European Conference on Computer Vision*, pages 446–463. Springer, 2022. [2](#)
- [39] Nibal Nayef, Yash Patel, Michal Busta, Pinaki Nath Chowdhury, Dimosthenis Karatzas, Wafa Khelif, Jiri Matas, Umada Pal, Jean-Christophe Burie, Cheng-lin Liu, et al. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition—rrc-mlt-2019. In *2019 International conference on document analysis and recognition (ICDAR)*, pages 1582–1587. IEEE, 2019. [5](#)
- [40] Gaurav Patel, Jan P Allebach, and Qiang Qiu. Seq-ups: Sequential uncertainty-aware pseudo-label selection for semi-supervised text recognition. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 6180–6190, 2023. [2](#), [7](#)
- [41] Yash Patel and Jiří Matas. Feds-filtered edit distance surrogate. In *International Conference on Document Analysis and Recognition*, pages 171–186. Springer, 2021. [1](#), [7](#)
- [42] Hieu Pham, Zihang Dai, Qizhe Xie, Minh-Thang Luong, and Quoc V. Le. Meta pseudo labels, 2021. [2](#)
- [43] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *Proceedings of the IEEE international conference on computer vision*, pages 569–576, 2013. [2](#), [5](#)
- [44] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications*, 41(18):8027–8048, 2014. [2](#), [5](#)
- [45] Mamshad Nayeem Rizve, Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. In defense of pseudo-labeling: An uncertainty-aware pseudo-label selection framework for semi-supervised learning. *arXiv preprint arXiv:2101.06329*, 2021. [2](#)
- [46] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(11):2298–2304, 2016. [2](#), [6](#)
- [47] Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Robust scene text recognition with automatic rectification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4168–4176, 2016. [2](#)
- [48] 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](#)
- [49] Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. Icdar2017 competition on reading chinese text in the wild (rctw-17). In *2017 14th iapr international conference on document analysis and recognition (ICDAR)*, volume 1, pages 1429–1434. IEEE, 2017. [5](#)
- [50] Amanpreet Singh, Guan Pang, Mandy Toh, Jing Huang, Wojciech Galuba, and Tal Hassner. Textocr: Towards large-scale

- end-to-end reasoning for arbitrary-shaped scene text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8802–8812, 2021. 5
- [51] Leslie N Smith and Nicholay Topin. Super-convergence: Very fast training of neural networks using large learning rates. In *Artificial intelligence and machine learning for multi-domain operations applications*, volume 11006, pages 369–386. SPIE, 2019. 6
- [52] Qi Song, Qianyi Jiang, Lei Wang, Lingling Zhao, and Rui Zhang. Mugs: A multiple granularity semi-supervised method for text recognition. In *International Conference on Document Analysis and Recognition*, pages 173–188. Springer, 2023. 1
- [53] Cheng Sun, Juntao Cheng, and Cheng Du. Semi-and self-supervised learning for scene text recognition with fewer labels. In *Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, pages 295–307. Springer, 2022. 1
- [54] Yipeng Sun, Zihan Ni, Chee-Kheng Chng, Yuliang Liu, Canjie Luo, Chun Chet Ng, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, et al. Icdar 2019 competition on large-scale street view text with partial labeling-rrc-lsvt. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1557–1562. IEEE, 2019. 5
- [55] Hung Tran Tien and Thanh Duc Ngo. Unsupervised domain adaptation with imbalanced character distribution for scene text recognition. In *2023 IEEE International Conference on Image Processing (ICIP)*, pages 3493–3497. IEEE, 2023. 2
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 2
- [57] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016. 2, 5
- [58] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *2011 International conference on computer vision*, pages 1457–1464. IEEE, 2011. 2, 5
- [59] Peng Wang, Cheng Da, and Cong Yao. Multi-granularity prediction for scene text recognition. In *European Conference on Computer Vision*, pages 339–355. Springer, 2022. 2
- [60] Tianwei Wang, Yuanzhi Zhu, Lianwen Jin, Canjie Luo, Xiaoxue Chen, Yaqiang Wu, Qianying Wang, and Mingxiang Cai. Decoupled attention network for text recognition. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 12216–12224, 2020. 2
- [61] Thomas Westfechtel, Hao-Wei Yeh, Dexuan Zhang, and Tatsuya Harada. Gradual source domain expansion for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 1946–1955, 2024. 2
- [62] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V Le. Self-training with noisy student improves imagenet classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10687–10698, 2020. 2
- [63] Xiaoyu Yue, Zhanghui Kuang, Chenhao Lin, Hongbin Sun, and Wayne Zhang. Robustscanner: Dynamically enhancing positional clues for robust text recognition. In *European Conference on Computer Vision*, pages 135–151. Springer, 2020. 2
- [64] Fangneng Zhan, Chuhui Xue, and Shijian Lu. Ga-dan: Geometry-aware domain adaptation network for scene text detection and recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9105–9115, 2019. 2
- [65] Boqiang Zhang, Hongtao Xie, Yuxin Wang, Jianjun Xu, and Yongdong Zhang. Linguistic more: Taking a further step toward efficient and accurate scene text recognition. *arXiv preprint arXiv:2305.05140*, 2023. 2
- [66] Rui Zhang, Yongsheng Zhou, Qianyi Jiang, Qi Song, Nan Li, Kai Zhou, Lei Wang, Dong Wang, Minghui Liao, Mingkun Yang, et al. Icdar 2019 robust reading challenge on reading chinese text on signboard. In *2019 international conference on document analysis and recognition (ICDAR)*, pages 1577–1581. IEEE, 2019. 2, 5
- [67] Ying Zhang, Lionel Gueguen, Ilya Zharkov, Peter Zhang, Keith Seifert, and Ben Kadlec. Uber-text: A large-scale dataset for optical character recognition from street-level imagery. In *SUNw: Scene Understanding Workshop-CVPR*, volume 2017, page 5, 2017. 2, 5
- [68] Yaping Zhang, Shuai Nie, Shan Liang, and Wenju Liu. Robust text image recognition via adversarial sequence-to-sequence domain adaptation. *IEEE Transactions on Image Processing*, 30:3922–3933, 2021. 2
- [69] Yaping Zhang, Shuai Nie, Wenju Liu, Xing Xu, Dongxiang Zhang, and Heng Tao Shen. Sequence-to-sequence domain adaptation network for robust text image recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2740–2749, 2019. 2
- [70] Shuai Zhao, Ruijie Quan, Linchao Zhu, and Yi Yang. Clip4str: A simple baseline for scene text recognition with pre-trained vision-language model. *arXiv preprint arXiv:2305.14014*, 2023. 2
- [71] Caiyuan Zheng, Hui Li, Seon-Min Rhee, Seungju Han, Jae-Joon Han, and Peng Wang. Pushing the performance limit of scene text recognizer without human annotation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14116–14125, 2022. 7
- [72] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 4