

Out of Length Text Recognition with Sub-String Matching

Yongkun Du¹, Zhineng Chen^{1*}, Caiyan Jia², Xieping Gao³, Yu-Gang Jiang¹

¹School of Computer Science, Fudan University, China

²School of Computer Science and Technology, Beijing Jiaotong University, China

³Laboratory for Artificial Intelligence and International Communication, Hunan Normal University, China
ykdu23@m.fudan.edu.cn, {zhinchen, ygi}@fudan.edu.cn, cyjia@bjtu.edu.cn, xpgao@hunnu.edu.cn

Abstract

Scene Text Recognition (STR) methods have demonstrated robust performance in word-level text recognition. However, in real applications the text image is sometimes long due to detected with multiple horizontal words. It triggers the requirement to build long text recognition models from readily available short (i.e., word-level) text datasets, which has been less studied previously. In this paper, we term this task Out of Length (OOL) text recognition. We establish the first Long Text Benchmark (LTB) to facilitate the assessment of different methods in long text recognition. Meanwhile, we propose a novel method called OOL Text Recognition with sub-String Matching (SMTR). SMTR comprises two cross-attention-based modules: one encodes a sub-string containing multiple characters into next and previous queries, and the other employs the queries to attend to the image features, matching the sub-string and simultaneously recognizing its next and previous character. SMTR can recognize text of arbitrary length by iterating the process above. To avoid being trapped in recognizing highly similar sub-strings, we introduce a regularization training to compel SMTR to effectively discover subtle differences between similar sub-strings for precise matching. In addition, we propose an inference augmentation strategy to alleviate confusion caused by identical sub-strings in the same text and improve the overall recognition efficiency. Extensive experimental results reveal that SMTR, even when trained exclusively on short text, outperforms existing methods in public short text benchmarks and exhibits a clear advantage on LTB.

Code — <https://github.com/Topdu/OpenOCR>

1 Introduction

Extracting text from natural images is a crucial and well-established task, typically encompassing scene text detection and recognition. In scene text recognition (STR), past research has achieved significant progress in word-level text recognition. However, in real-world applications, the text is not always detected as individual words, but a line-level text sometimes. That is, the text instance detected by scene text detectors may contain multiple horizontal words, presenting a long text recognition challenge. It is worth noting that

*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

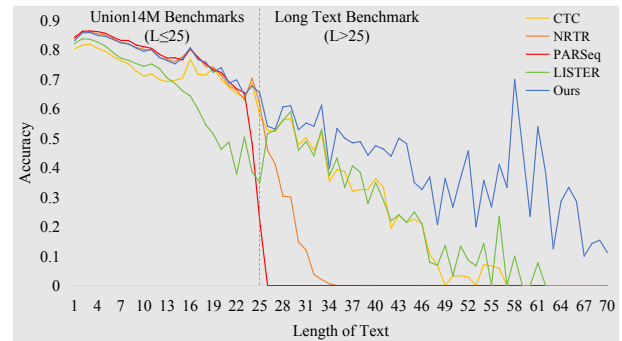


Figure 1: Attention-based methods like NRTR (Sheng, Chen, and Xu 2019) and PARSeq (Bautista and Atienza 2022) perform well in short text recognition, while the CTC-based FocalSVTR (CTC) and the length-insensitive LISTER (Cheng et al. 2023) outperform them in long text but worse in short text. Our proposed SMTR gets superior performance in both short and long text recognition.

existing STR datasets are predominantly comprised of word-level text, where long text like the aforementioned is scarce. Therefore, accurately recognizing long text by utilizing solely short (i.e., word-level) datasets has emerged as a promising yet challenging frontier in STR, which we term Out of Length (OOL) text recognition in this paper.

Existing STR models are mostly designed towards recognizing short text with no more than 25 characters, i.e., $L \leq 25$. We establish the first Long Text Benchmark (LTB) that focuses on the long text ($L > 25$). As shown in Fig. 1, we evaluate several popular STR models (Sheng, Chen, and Xu 2019; Du et al. 2022; Bautista and Atienza 2022; Cheng et al. 2023) on LTB. The results show that the CTC-based method (Du et al. 2022) and LISTER (Cheng et al. 2023) can recognize long text due to their length extrapolation capabilities. However, their accuracy declines rapidly as text length increases. Additionally, these methods show a noticeable performance gap in short text recognition compared to the competitors (Sheng, Chen, and Xu 2019; Bautista and Atienza 2022), due to not being equipped with an advanced decoder. On the other hand, NRTR (Sheng, Chen, and Xu 2019) and PARSeq (Bautista and Atienza 2022), as representatives of the attention-based methods, employ the

self-attention mechanism to encode the decoded characters along with absolute positional information (Vaswani et al. 2017). It serves as the linguistic context or position-involved representation to aid decoders in accurately recognizing characters. This powerful approach enables these methods to achieve state-of-the-art results in short text recognition. However, they are unable to effectively learn the representation beyond the length of training text, which limits them from processing long text when trained solely on short text. Similar limitations are also observed in other attention-based methods (Shi et al. 2019; Wang et al. 2021; Fang et al. 2021; Du et al. 2023; Zhang et al. 2023; Zheng et al. 2024).

In this paper, we introduce a novel method termed OOL Text Recognition with sub-String Matching (SMTR), which achieves OOL text recognition by innovatively leveraging string-matching techniques. Specifically, SMTR recognizes text by first matching a specified sub-string within an image, and then positioning and recognizing the Next and Previous characters of the sub-string. SMTR can recognize text of arbitrary length by iterating the process above. It gets rid of the absolute position and recognizes text fully relying on sub-string identification and extrapolation. This new paradigm is reasonable as both short and long text can be decomposed into a series of sub-string units. The matching process only compares a portion of the text image with the sub-string, regardless of whether the whole text is short or long. Based on this fact, SMTR learns to match sub-strings by using only short text images but can still recognize long text.

To implement this, SMTR develops a lightweight architecture with a sub-string encoder and a sub-string matcher. The former encodes a sub-string as two representations referred to as next and previous queries. While the latter directs the two queries to attend to visual features, aiming to accurately match the sub-string and recognize the next and previous characters. In this new recognition paradigm, similar or repeated sub-strings in one text image inevitably hinder precise sub-string matching. To effectively screen for similar sub-strings, we propose a regularization training strategy that encourages SMTR to pay attention to the subtle differences between sub-strings, thus better distinguishing them. In addition, we introduce an inference augmentation strategy which breaks a long text image into multiple sub-images for independent processing, thus significantly alleviating the side affection caused by repeated sub-strings. Moreover, by processing these sub-images in parallel, SMTR improves the overall recognition efficiency remarkably. Experimental results demonstrate that SMTR achieves highly competitive performance on challenging short text benchmarks and exhibits a noteworthy advantage on LTB. The contributions are summarised as follows:

- We, for the first time, term the requirement of building long text recognition models based on short text datasets, as OOL text recognition challenge, and establish a long text benchmark called LTB to assess the long text recognition capability of STR models.
- We propose SMTR, a novel method that elegantly incorporates string-matching techniques to address the OOL challenge. Meanwhile, regularization training and infer-

ence augmentation strategies are proposed to ensure precise recognition following this pipeline.

- SMTR achieves state-of-the-art performance on both public short text benchmark and LTB, utilizing merely short text for training. Our study enriches the STR schemes in handling diverse real-world applications.

2 Related Work

Existing STR methods mainly focus on word-level text recognition. Among them, attention-based (Lee and Osindero 2016) methods are intensively studied for their impressive performance. Some of these methods (Bautista and Atienza 2022; Yue et al. 2020; Wang et al. 2021; Du et al. 2023; Fang et al. 2021; Zhang et al. 2023; Fang et al. 2023; Zhao et al. 2024; Du et al. 2024a) use learnable position embeddings to accurately learn context information. However, because learnable position embeddings are typically not scalable, these methods can only recognize short text. In contrast, methods (Shi et al. 2019; Li et al. 2019; Sheng, Chen, and Xu 2019; Qiao et al. 2021; Zheng et al. 2024; Zhang et al. 2024a) employ LSTM (Hochreiter and Schmidhuber 1997; Chung et al. 2014) or masked self-attention mechanism with sinusoidal positional encoding (Lee and Osindero 2016) to model context. They have some length extrapolation capability. Nevertheless, these models are primarily trained on short text. Therefore their inference ability is limited when dealing with long text, resulting in a significant decrease in recognition accuracy. Additionally, some methods do not rely on well-designed attention-based decoders, such as CTC-based (Graves et al. 2006) methods (Shi, Bai, and Yao 2017; Hu et al. 2020; Du et al. 2022) and length-insensitive LISTER (Cheng et al. 2023). They do not perform well on short text recognition although exhibiting slightly stronger capability in long text recognition.

SMTR circumvents the above problems and recognizes text by sub-string matching. This scheme ensures SMTR’s applicability to long text even when only seen short text images, as sub-string is an essential component for both short and long text.

3 Method

Fig. 2 illustrates the overall architecture of SMTR. Given a text image $X \in \mathbb{R}^{3 \times H \times W}$, the image encoder extracts its image embeddings E_I . A sub-string (S) is encoded into both next and previous hidden representation (Q_n and Q_p) by the sub-string encoder. Q_n and Q_p are then fed into the sub-string matcher, which implicitly matches the sub-string in the text image and recognizes its next and previous characters bidirectionally. In addition, we discuss issues raised within this new string-matching-based paradigm and propose our solutions.

3.1 Image Encoder

To accommodate various aspect ratios of the input X , we develop a dedicated image encoder termed FocalSVTR to extract visual features as in LISTER (Cheng et al. 2023). Firstly, the input X is divided into patch embeddings (\in

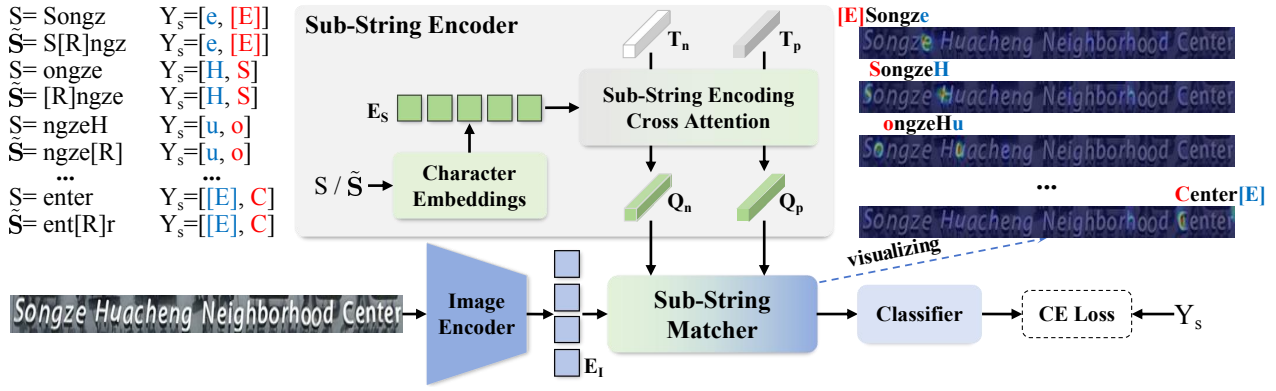


Figure 2: Overview of SMTR. SMTR is a lightweight recognizer consisting of two cross-attention-based modules, i.e., sub-string encoder and sub-string matcher. S and \tilde{S} denote a sub-string and its regularized counterpart. Y_s denotes the next-previous character set of the sub-string. [R] denotes a randomly selected character from the character set.

$\mathbb{R}^{C_0 \times \frac{H}{4} \times \frac{W}{4}}$) by two convolution with stride 2. Then, following SVTR (Du et al. 2022), three stages comprising [6, 6, 6] layers of focal modulation (Yang et al. 2022) are stacked. At the end of the first and second stages, C_0 is extended with a factor of 2 by a convolution. In particular, at the end of the second stage, the convolution downsamples the height of the patch embeddings to $\frac{H}{8}$. Finally, the output features ($\in \mathbb{R}^{C \times \frac{H}{8} \times \frac{W}{4}}$) are flattened to obtain image embeddings $E_I \in \mathbb{R}^{C \times \frac{HW}{32}}$, where $C = 4C_0$.

3.2 Sub-String Encoder

Sub-String and Label Generation. For a labeled text image, all sub-strings S and their next-previous character sets $Y_s = [Y_n, Y_p]$ can be generated as follows. Taking “datours” in Fig. 4 as an example. It is first prefixed and suffixed with “[B] $_{l_s}$ ”, resulting in the formatted label: “[B] $_{l_s}$ datours[B] $_{l_s}$ ”. Here, “[B] $_{l_s}$ ” represents a string with l_s blank characters [B], and l_s is the predefined sub-string length. After that, the formatted label is traversed by a window of size l_s to obtain all sub-strings and their Y_s labels. If a sub-string contains [B] on one side, the corresponding label in Y_s is replaced with [P], indicating that there is no additional character on that side.

Fig. 4 shows all sub-strings of “datours” and their corresponding Y_s when l_s is set to 5. The sub-string is initialized with “[B] $_5$ ” and $Y_s = [d, s]$, indicating the start of the next and previous inferences. For a text of length L , the number of sub-strings is $2L + 1$ when $L < l_s$, and $L + l_s$ otherwise.

Sub-String Encoding. Given a sub-string S , firstly, each character in S is converted into a C -dim embedding via a character embedding layer, and the sub-string embedding $E_s \in \mathbb{R}^{l_s \times C}$ is obtained by feature concatenation. Then, the next token $T_n \in \mathbb{R}^{1 \times C}$ is introduced, which is a shared and learnable token. T_n acts as a query and performs cross-attention with E_s , obtaining the next sub-string hidden representation $Q_n \in \mathbb{R}^{1 \times C}$. Similarly, the previous token $T_p \in \mathbb{R}^{1 \times C}$ is introduced and generates the previous sub-string hidden representation $Q_p \in \mathbb{R}^{1 \times C}$. Assuming there are h attention heads, each with different param-

eters $W_i^q, W_i^k, W_i^v \in \mathbb{R}^{C \times C_h}$, i denotes the index of the attention head. The above process can be formulated as:

$$Q_z = \text{MHead}(T_z, E_s, E_s) + T_z$$

$$\text{MHead}(Q, K, V) = \text{Concat}(\text{head}_1, \text{head}_2, \dots, \text{head}_h)$$

$$\text{head}_i = A_i(VW_i^v)$$

$$A_i = \sigma \left((QW_i^q)(KW_i^k) C_h^{-0.5} \right) \in \mathbb{R}^{1 \times \frac{HW}{32}} \quad (1)$$

where z is n or p , $C_h = \frac{C}{h}$, and σ is the Softmax function.

3.3 Sub-String Matcher

We use one cross-attention (Vaswani et al. 2017) to implement sub-string matching, where the sub-string hidden representation Q_n or Q_p as query and image embeddings E_I as key and value. The process can be expressed as:

$$\text{Matcher}(Q_z, E_I) = \text{MHead}(Q_z, E_I, E_I) \quad (2)$$

Taking Q_n as an example, the sub-string matcher computes the attention map A_i , which focuses on the position of the next character of the sub-string, as shown in the attention maps in Fig. 2. Consequently, the output of the sub-string matcher represents the next character feature $F_n = \text{Matcher}(Q_n, E_I) \in \mathbb{R}^{1 \times C}$, which then undergoes a Classifier ($\in \mathbb{R}^{C \times (N_c + 1)}$) to generate prediction $\tilde{Y}_n \in \mathbb{R}^{1 \times (N_c + 1)}$, i.e., $\tilde{Y}_n = \text{Classifier}(F_n)$, where $+1$ is for the end symbol [E]. In the same way, the previous character prediction is obtained by $\tilde{Y}_p = \text{Classifier}(F_p)$, $F_p = \text{Matcher}(Q_p, E_I)$.

3.4 Optimization Objective

During training, for a text instance with N sub-strings, SMTR predicts the next and previous characters (\tilde{Y}_n and \tilde{Y}_p) for each sub-string. The loss \mathcal{L} is computed by comparing the label $Y_s = [Y_n, Y_p]$ with $\tilde{Y}_s = [\tilde{Y}_n, \tilde{Y}_p]$:

$$\mathcal{L} = \frac{1}{K} \sum_{i=1}^N (ce(\sigma(\tilde{Y}_n^i), Y_n^i) + ce(\sigma(\tilde{Y}_p^i), Y_p^i)) \quad (3)$$

$$ce(\tilde{y}, y) = \begin{cases} -\sum_{c=1}^{N_c} y_c \log(\tilde{y}_c), & y \neq [\text{P}]. \\ 0, & y = [\text{P}]. \end{cases}$$

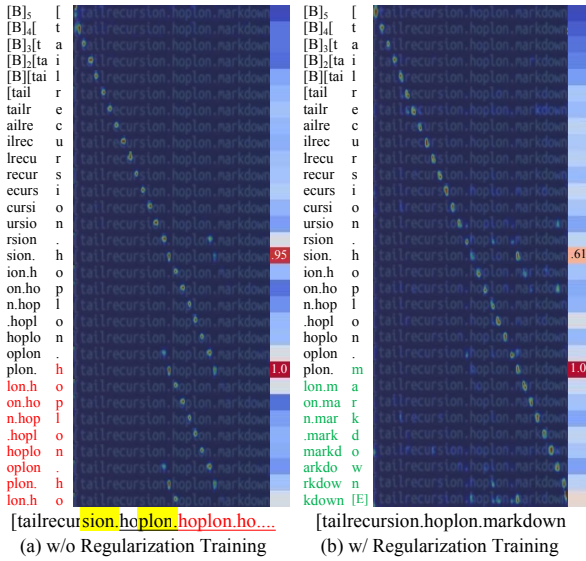


Figure 3: Attention maps of SMTR w/o (left) and w/ (right) the proposed regularization training, which rectifies recognition errors (red) caused by similar sub-strings (yellow).

where label [P] is not involved in the loss computation, K is the number of valid loss terms. It is equal to $2N - 2L$ when $L < l_s$, and $2N - 2(l_s - 1)$ otherwise.

3.5 Regularization Training

SMTR performs matching and recognition by aligning the sub-string embedding within the image feature space. However, this process can be compromised by similar or repeated sub-strings, as they provide quite similar visual features and context. We categorize this problem as similar distraction and repeated confusion, and propose two strategies to alleviate them. The first is regularization training for similar distraction described as follows.

SMTR is easily disturbed by sub-strings with similar visual appearance and generates error recognition. As shown in Fig. 3(a), the next character after “plon.” is mistakenly identified as the next character of “sion.”, resulting in an undesired circular recognition. To overcome this issue, SMTR should be endowed with the ability to discover subtle differences between similar sub-strings. To this end, we propose regularization training (RT), which alleviates this problem by generating similar sub-strings \tilde{S} of S for training reinforcement. As depicted in the top-left side of Fig. 2, one character in S is replaced with [R], a character randomly selected from the character set, to obtain \tilde{S} , and keep its label Y_s unchanged. By doing this, the number of similar sub-strings is largely enriched during training. SMTR is compelled to leverage all the l_s characters, rather than a few nearby ones, to distinguish similar sub-strings appearing in the same text image, which therefore can be better identified.

3.6 Inference Augmentation

We then introduce the inference augmentation (IA) proposed to mitigate the repeated confusion as follows.




Figure 4: Illustration of SMTR base inference process, where [E] means the end token for inference termination.

Algorithm 1: Base Inference Pseudo-code in Python

```
def Inference (Img, Tok, Ss=[0]*l_s, Se=[]):
    # Img: Input Image, 3 x H x W.
    # Tok: Next or Previous Token, 1 x C.
    # Ss/Se: The start/end sub-string.
    # [0]: The [B].
    # l_s: The length of sub-string.
    Result = [], S = Ss
    IE = ImageEncoder (Img) # HW/32 x C
    While True:
        Q = SubStringEncoding (Tok, S)
        O = SubStringMatcher (Q, IE)
        Char = Classifier (O).argmax (-1)
        if Tok is T_n:
            S = S[1:] + [Char]
        else:
            S = [Char] + S[:1]
        # EOS: The index of end symbol.
        if Char == EOS or S == Se:
            break
        Result.append (Char)
    return Result
    # Result is the recognition result.
```

We first introduce the base inference of SMTR. Since SMTR does not know which sub-strings are contained in the text image, it implements the decoding inference with the sub-strings beginning with $[B]_{l_s}$, and iteratively predicts the next and previous characters as shown in Fig. 4. The process is terminated when the end token [E] is predicted. Alg. 1 provides a detailed description of this process.

The base inference fails to distinguish repeated sub-strings in one image, such as the SMTR w/o IA result in Fig. 5. Tab. 1 shows the statistics of repeated sub-strings in typical short text datasets and LTB. The high percentage of repeated sub-strings (near 10%) in LTB seriously affects the recognition. Consequently, we propose IA that first slices the long text image into three short sub-images and then employs a split-merge scheme for recognition. This process is depicted in Fig. 5. Specifically, the text image is sliced into three sub-images whose width is halved, i.e., the left sub-image Img_1 , the right sub-image Img_2 and the central sub-image Img_3 . Note that both Img_1 and Img_2 are overlapped with Img_3 . Then, by taking advantage of the bidirectional recognition property of SMTR, we recognize Img_1 and Img_2 following the next and previous prediction paths, respectively, i.e., $Inference (Img_1, T_n)$ and $Inference (Img_2, T_p)$. The process obtains the results shown in the top line of Fig. 5 in roughly half of the

Dataset	Samples	Repeat S	
Union14M-L Training	3,230,742	885	0.03%
Union14M-Benchmark	409,383	22	0.01%
Common Benchmarks	7,248	0	0.00%
LTB w/o IA	4,789	472	9.86%
LTB w/ IA	4,789	20	0.42%

Table 1: The percentage of text instances with repeated sub-strings ($l_s \geq 5$) in different datasets. Inference augmentation (IA) drastically reduces the percentage in LTB.

inference iterations, as both predictions can be carried out in parallel. Since the slice may operate on characters and lead to mis-recognition, we pick out recognized sub-strings also appearing in Img_3 from both sides according to attention maps (highlighted by yellow in Fig. 5), i.e., S_s and S_e . In the following, we predict Img_3 with $\text{Inference}(\text{Img}_3, T_n, S_s, S_e)$, which gets the central characters. Finally, the full result, e.g., the *SMTR w/ IA* result in Fig. 5, is obtained by a simple post-processing.

It is seen in Tab. 1 that IA drastically reduces the percentage of repeated sub-strings in LTB from nearly 10% to 0.4%, thus significantly alleviating the repeated confusion. The effectiveness of IA will be demonstrated in Tab. 3.

3.7 Long Text Benchmark

Existing public datasets (Jung et al. 2011; Nagy, Dicker, and Meyer-Wegener 2012; Veit et al. 2016; Shi et al. 2017; Zhang et al. 2017; Liu et al. 2017; Mathew, Jain, and Jawahar 2017; He et al. 2018; Nayef et al. 2019; Sun et al. 2019; Chng et al. 2019; Zhang et al. 2019; Krylov, Nosov, and Sovrasov 2021; Singh et al. 2021; Long et al. 2022) predominantly consist of short text ($L \leq 25$). However, they also have a few instances with more than 25 characters. These instances are usually regarded as noise and discarded in developing traditional short text models. We collect these instances and construct a Long Text Benchmark (*LTB*). *LTB* contains nearly 4.8k samples whose length is above 25. Since all samples are excluded from the training process, *LTB* establishes a new benchmark exclusively for evaluating the performance of STR models on long text recognition. To assess the impact of length variation on recognition, we divide *LTB* into three parts based on text length: [26, 35], [36, 55], and ≥ 56 , as shown in Tab. 3.

4 Experiments

4.1 Datasets and Implementation Details

We evaluate SMTR on both English and Chinese datasets. For English, our models are trained on Union14M-L (Jiang et al. 2023), which contains about 3.2 million real-world text images ($L \leq 25$). Then, the models are tested on both LTB and two short text benchmarks: (1) Common benchmarks, i.e., ICDAR 2013 (*IC13*) (KaratzasAU et al. 2013), Street View Text (*SVT*) (Wang, Babenko, and Belongie 2011), IIIT5K-Words (*IIIT*) (Mishra, Karteek, and Jawahar 2012), ICDAR 2015 (*IC15*) (Karatzas et al. 2015), Street View Text-Perspective (*SVTP*) (Phan et al. 2013) and CUTE80 (*CUTE*) (Anhar et al. 2014). For *IC13* and




Figure 5: Illustration of inference augmentation (IA) for long text recognition. Text spaces are manually inserted.

IC15, we use the versions with 857 and 1,811 images, respectively. (2) Union14M-Benchmark, which includes seven challenging subsets: Curve (*CUR*), Multi-Oriented (*MLO*), Artistic (*ART*), Contextless (*CTL*), Salient (*SAL*), Multi-Words (*MLW*) and General (*GEN*). Note that the 6,600 samples overlapping between Union14M-L and Union14M-Benchmark are filtered (Du et al. 2024b). For Chinese, we use Chinese text recognition (CTR) dataset (Chen et al. 2021), which contains four subsets: *Scene*, *Web*, *Document* (*Doc*) and *Hand-Writing* (*HW*). We train the model on the whole training set and use *Scene* validation subset to determine the best model, which is assessed on the test subsets.

We use AdamW optimizer (Loshchilov and Hutter 2019) with a weight decay of 0.05 for training. The LR is set to 6.5×10^{-4} and batchsize is set to 1024. One cycle LR scheduler (I. Loshchilov and Hutter 2017) with 1.5/4.5 epochs linear warm-up is used in all the 20/100 epochs, where a/b means a for English and b for Chinese. Regarding the aspect ratio, all images are resized to a maximum pixel size of 32×128 if the aspect ratio is less than 4, otherwise, it is resized to $H = 32$ and W up to 384. Word accuracy is used as the evaluation metric. Data augmentation like rotation, perspective distortion, motion blur and gaussian noise, are randomly performed and the maximum text length is set to 25 during training. The size of the character set N_c is set to 96 for English and 6625 (Li et al. 2022) for Chinese. The number of attention heads h in sub-string Encoder and Matcher are set to $\frac{C}{32}$ and 2, respectively. All models are trained with mixed-precision on 4 RTX 3090 GPUs.

4.2 Ablation Study

Effectiveness of Regularization Training (RT). As shown in Fig. 3(b), RT mitigates the sub-string mismatching effectively. With RT, SMTR successfully discriminates two similar sub-strings (“sion.” and “plon.”) and recognizes the text accurately. Tab. 2 quantifies the effectiveness of RT. It takes effects for both short and long text, where a notable 12.07% accuracy improvement is observed on LTB. The result indicates similar sub-strings are much better identified.

Tab. 2 also ablates l_s and \tilde{S} , the number of regularized sub-strings. The results indicate that $l_s > 5$ is not necessary. In addition, increasing \tilde{S} enhances the accuracy on both short and long text, again validating the positive effect of RT. Therefore, SMTR selects $l_s = 5$ and $\tilde{S} = 2$.

In addition, to qualitatively assess RT in differentiating similar sub-strings, we examine the cosine similarity between $Q_n^{plon.}$, the embedding of sub-string “sion.”, and oth-

\tilde{S}	l_s	Common	U14M	LTB	Avg
w/o	5	95.65	83.58	32.30	70.51
1	5	95.94	84.14	44.37	74.82
	4	96.04	84.51	45.24	75.26
	5	95.90	85.00	47.00	75.97
2	6	95.64	84.26	45.06	74.99
	7	95.77	83.80	40.55	73.37

Table 2: Ablation study on regularization training and l_s .

Method	$L_{[26,35]}$ 3376	$L_{[36,55]}$ 1147	$L_{\geq 56}$ 266	W-Avg	A-Avg
FocalSVTR	51.04	25.37	0.38	42.08	25.59
AR-STR	25.53	0.00	0.00	18.00	8.51
PARSeq [†] (2022)	0.00	0.00	0.00	0.00	0.00
LISTER [†] (2023)	51.16	26.59	2.26	42.56	26.67
SMTR w/o IA	53.06	37.05	13.16	47.00	34.42
SMTR w/ IA	55.48	43.68	25.56	50.99	41.57

Table 3: Model comparing results on LTB.

ers. As shown in the rightmost column in Fig. 3(a), $Q_n^{plon.}$ has a similarity score of 1.0 with itself. Without RT, the similarity between $Q_n^{plon.}$ and $Q_n^{sion.}$, the embedding of its similar sub-string “sion.”, is as high as 0.95, which leads to SMTR not being able to distinguish between them well. In contrast, RT effectively reduces this similarity from 0.95 to 0.61. As shown in the attention map in Fig. 3(b), this similarity allows SMTR to avoid this distraction and correctly focus on the next character. This example demonstrates that RT successfully pushes away the representations of two similar sub-strings, thereby enhancing SMTR’s ability to distinguish between them and improving recognition accuracy.

Effectiveness of Inference Augmentation (IA). As shown in Tab. 3, IA largely improves the accuracy of SMTR on LTB, with weighted (W-Avg) and arithmetic (A-Avg) average increasing by 3.99% and 7.17%, respectively. Notably, accuracy gains from IA become more distinct as the text length increases, where 2.42%, 6.63%, and 12.4% improvements are observed across the three subsets. This trend arises because longer texts are more likely to contain repeating sub-strings. IA effectively eliminates these repetitions, as shown in Tab. 1, leading to substantial improvements in recognition accuracy.

Fair comparison with aligned methods on LTB. As aforementioned, some STR models, despite being designed for short text recognition, are also capable of recognizing long text, e.g., CTC-based and auto-regressive (AR) decoding ones. We align four of them in Tab. 3 with SMTR for comparison. Specifically, we uniformly adopt FocalSVTR as the image encoder, keeping their decoders unchanged, and then training them using the same hyper-parameters detailed in section 3.1. The results are presented in Tab. 3 and Tab. 4, where FocalSVTR is the encoder combining with the naive CTC decoder, AR-STR is the encoder plus with a standard AR decode composed of two transformer units (Vaswani et al. 2017), and [†] denotes our aligned reproductions.

As seen in Tab. 3, SMTR consistently outperforms the compared models in terms of accuracy, and their gaps are

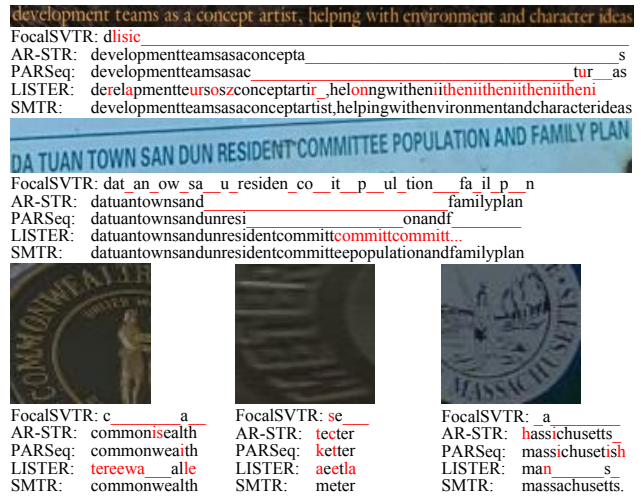


Figure 6: Illustration of the predictions of two long text images from LTB and three short text images from Union14M-Benchmark. In addition, their attention maps are visualized in Supplementary for qualitative analysis.

sharply enlarged as the text length increases. When inspecting the arithmetic average for $L > 35$, SMTR outperforms FocalSVTR, AR-STR and LISTER by 15.98%, 33.06%, and 14.90%, respectively. The first two long text examples in Fig. 6 show that all four comparing methods suffer from the problem of missing characters. In addition, AR-STR and PARSeq are only able to recognize the front characters due to employing absolute positional encoding and only trained on short text. LISTER encounters the problem of circular recognition due to capturing incorrect neighbor characters. In contrast, SMTR effectively recognizes these instances by sub-string matching. These results convincingly verify the superiority of SMTR in recognizing long text and addressing the OOL challenge.

4.3 Comparison with State-of-the-arts

Results on English Benchmarks. We apply SMTR to Common benchmarks, Union14M-Benchmark and LTB, and compare it with existing STR models. The results are in Tab. 4. Compared with PARSeq, one leading method in short text benchmarks, SMTR exhibits quite similar accuracy. In addition, it correctly recognizes half of LTB text, which is incapable for PARSeq. Note that AR-STR also achieves very competitive results on short text benchmarks. However, its weak extrapolation capability restricts the performance on LTB. On the other hand, FocalSVTR performs decently on short text benchmarks, but on LTB it performs on par with LISTER, which is a dedicated model for long text recognition. The results above are in line with the observation in Fig. 1. To sum, by introducing a novel sub-string matching-based recognition paradigm, SMTR well recognizes both short and long text, which would be a promising property for many STR-related applications.

In the bottom of Fig. 6 we also present the recognition of three challenging short text images. CTC-based FocalSVTR

Method	Common Benchmarks							Union14M-Benchmark							LTB	Params ($\times 10^6$)	
	IC13	SVT	IIIT	IC15	SVTP	CUTE	Avg	CUR	MLO	ART	CTL	SAL	MLW	GEN			Avg
CRNN (2017)	91.8	83.8	90.8	71.8	70.4	80.9	81.58	19.4	4.5	34.2	44.0	16.7	35.7	60.4	30.70	-	8.3
SVTR-B* (2022)	97.5	96.4	97.8	89.3	91.0	96.2	94.72	85.4	87.4	68.9	79.5	84.3	79.1	81.8	80.91	-	24.6
FocalSVTR	97.3	96.3	98.2	87.4	88.4	96.2	93.97	77.7	62.4	65.7	78.6	71.6	81.3	79.2	73.80	42.1	14.7
DAN (2020)	95.2	88.6	95.5	78.3	79.9	86.1	87.26	46.0	22.8	49.3	61.6	44.6	61.2	67.0	50.44	0.0	27.7
SRN (2020)	94.7	89.5	95.5	79.1	83.9	91.3	89.00	49.7	20.0	50.7	61.0	43.9	51.5	62.7	48.50	0.0	54.7
RoScanner (2020)	95.7	92.4	96.8	86.4	83.9	93.8	91.50	66.2	54.2	61.4	72.7	60.1	74.2	75.7	66.36	0.0	48.0
ABINet (2021)	97.2	95.7	97.2	87.6	92.1	94.4	94.03	75.0	61.5	65.3	71.1	72.9	59.1	79.4	69.19	0.0	36.7
VisionLAN (2021)	95.1	91.3	96.3	83.6	85.4	92.4	90.68	70.7	57.2	56.7	63.8	67.6	47.3	74.2	62.50	0.0	32.8
PARSeq* (2022)	98.4	98.1	98.9	90.1	94.3	98.6	96.40	87.6	88.8	76.5	83.4	84.4	84.3	84.9	84.26	0.0	23.8
MATRn (2022)	97.9	96.9	98.2	88.2	94.1	97.9	95.53	80.5	64.7	71.1	74.8	79.4	67.6	77.9	73.71	0.0	44.2
BUSNet* (2024)	97.8	98.1	98.3	80.2	95.3	96.5	96.06	83.0	82.3	70.8	77.9	78.8	71.2	82.6	78.10	0.0	32.0
OTE (2024)	98.0	98.0	98.1	89.1	95.5	97.6	96.10	83.1	82.8	73.5	73.7	79.7	70.3	82.2	77.90	0.0	25.2
ASTER (2019)	92.6	88.9	94.3	77.7	80.5	86.5	86.75	38.4	13.0	41.8	52.9	31.9	49.8	66.7	42.07	-	27.2
NRTR (2019)	96.9	94.0	96.2	80.9	84.8	92.0	90.80	49.3	40.6	54.3	69.6	42.9	75.5	75.2	58.20	-	31.7
SAR (2019)	96.0	92.4	96.6	82.0	85.7	92.7	90.90	68.9	56.9	60.6	73.3	60.1	74.6	76.0	67.20	-	57.7
MAERec (2023)	97.6	96.8	98.0	87.1	93.2	97.9	95.10	81.4	71.4	72.0	82.0	78.5	82.4	82.5	78.60	-	35.7
CDistNet* (2024)	97.8	97.1	98.7	89.6	93.5	96.9	95.59	81.7	77.1	72.6	78.2	79.9	79.7	81.1	78.62	-	65.5
LISTER* (2023)	97.4	98.1	98.2	89.2	93.5	95.5	95.33	71.6	55.9	68.9	76.4	68.1	80.2	80.9	71.72	42.6	49.9
AR-STR	98.1	98.3	99.0	89.8	94.2	97.9	96.22	88.1	91.2	76.2	83.3	83.2	85.6	84.9	84.62	18.0	19.3
SMTR	98.3	97.4	99.0	90.1	92.7	97.9	95.90	89.1	87.7	76.8	83.9	84.6	89.3	83.7	85.00	51.0	15.8

Table 4: Results on short text benchmarks and LTB when the models are trained on Union14M-L (2023) datasets. ”-” means the model cannot recognize long text directly. For short text benchmarks, we use the next inference. * denotes that the results on short benchmarks is obtained by training the model on Union14M-L using the code they released.

and LISTER show severe mis-recognition due to their reliance on a priori of regular text. In contrast, AR-STR and PARSeq can recognize a majority of the text, encountering errors only when dealing with particularly challenging characters. SMTR accurately identifies all the characters, showing robustness in handling challenging short text.

Results on Chinese Benchmarks. The OOL challenge is pervasive across languages. To verify the multilingual adaptability of SMTR, we conduct evaluations on the Chinese text recognition benchmark (CTR) (Chen et al. 2021). In line with CCR-CLIP (Yu et al. 2023), SMTR refrains from utilizing data augmentation (w/o Aug) to ensure a fair comparison. As presented in Tab. 5, SMTR outperforms LISTER by 6.83% on *Scene* subset. Note that SMTR gets this result without training on long text while LISTER has no imposed length limitation during training. Despite this discrepancy, SMTR still maintains an accuracy gain of 10.61% over LISTER in Chinese long text ($L_{>25}$). Meanwhile, SMTR achieves a new state-of-the-art on CTR, boosting an accuracy gain of 3.53% compared to CAM (Yang et al. 2024), the previous best method. These results demonstrate SMTR’s great adaptability in multilingual recognition.

5 Conclusion

In this paper, we point out that, in real applications, STR models sometimes have to recognize long text images and existing models trained on short text images cannot accomplish this task well. We term this the OOL text recognition challenge. We construct the LTB dataset for long text recognition assessment, and propose a novel SMTR that employs a sub-string-matching-based paradigm to overcome this challenge. SMTR implements the recognition by iteratively predicting the next and previous characters of a sub-

Method	Scene	Web	Doc	HW	Avg	$L_{>25}$	Params $\times 10^6$
CRNN (2017)	53.4	57.0	96.6	50.8	64.45	-	12.4
ASTER (2019)	61.3	51.7	96.2	37.0	61.55	-	27.2
MORAN (2019)	54.6	31.5	86.1	16.2	47.10	-	28.5
SAR (2019)	59.7	58.0	95.7	36.5	62.48	-	27.8
SEED (2020)	44.7	28.1	91.4	21.0	46.30	-	36.1
MASTER (2021)	62.8	52.1	84.4	26.9	56.55	-	62.8
ABINet (2021)	66.6	63.2	98.2	53.1	70.28	0.0	53.1
TransOCR (2021)	71.3	64.8	97.1	53.0	71.55	-	83.9
SVTR-B (2022)	71.7	73.8	98.2	52.2	73.98	-	26.3
CCR-CLIP (2023)	71.3	69.2	98.3	60.3	74.78	-	62.0
LISTER (2023)	73.0	-	-	-	-	35.00	55.0
DCTC (2024b)	73.9	68.5	99.4	51.0	73.20	-	40.8
CAM (2024)	76.0	69.3	98.1	59.2	76.80	-	135
SMTR w/o Aug	79.8	80.6	99.1	61.9	80.33	45.61	20.8
SMTR w/ Aug	83.4	83.0	99.3	65.1	82.68	49.41	20.8

Table 5: Results of different models on CTR dataset. $L_{>25}$ means long text recognition results on *Scene*.

string. It is capable of recognizing both short and long text. To make the recognition effective in this new paradigm, we introduce regularization training to suppress distractions caused by similar sub-strings, and inference augmentation to alleviate confusion caused by repeated sub-strings and improve recognition efficiency. Experimental results show that SMTR not only surpasses existing methods by a clear margin on LTB, but is also highly competitive on challenging short text benchmarks. Nevertheless, SMTR employs an iterative recognition paradigm in inference thus the speed is not fast. Hence, our future research will be devoted to addressing this problem.

Acknowledgments

This work was supported by National Key R&D Program of China (No. 2022YFB3104703) and in part by the National Natural Science Foundation of China (No. 32341012).

References

- Anhar, R.; Palaiahnakote, S.; Chan, C. S.; and Tan, C. L. 2014. A robust arbitrary text detection system for natural scene images. *Expert Syst. Appl.*, 41(18): 8027–8048.
- Bautista, D.; and Atienza, R. 2022. Scene Text Recognition with Permuted Autoregressive Sequence Models. In *ECCV*, 178–196.
- Chen, J.; Li, B.; and Xue, X. 2021. Scene Text Telescope: Text-Focused Scene Image Super-Resolution. In *CVPR*, 12021–12030.
- Chen, J.; Yu, H.; Ma, J.; Guan, M.; Xu, X.; Wang, X.; Qu, S.; Li, B.; and Xue, X. 2021. Benchmarking Chinese Text Recognition: Datasets, Baselines, and an Empirical Study. *CoRR*, abs/2112.15093.
- Cheng, C.; Wang, P.; Da, C.; Zheng, Q.; and Yao, C. 2023. LISTER: Neighbor Decoding for Length-Insensitive Scene Text Recognition. In *ICCV*, 19484–19494.
- Chng, C.; Liu, Y.; Sun, Y.; Ng, C.; Luo, C.; N.; Fang, C.; Zhang, S.; Han, J.; Ding, E.; et al. 2019. ICDAR2019 robust reading challenge on arbitrary-shaped text-rrc-art. In *ICDAR*, 1571–1576.
- Chung, J.; Gülçehre, Ç.; Cho, K.; and Bengio, Y. 2014. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. *CoRR*, abs/1412.3555.
- Du, Y.; Chen, Z.; Jia, C.; Yin, X.; Li, C.; Du, Y.; and Jiang, Y. 2023. Context Perception Parallel Decoder for Scene Text Recognition. *CoRR*, abs/2307.12270.
- Du, Y.; Chen, Z.; Jia, C.; Yin, X.; Zheng, T.; Li, C.; Du, Y.; and Jiang, Y. 2022. SVTR: Scene Text Recognition with a Single Visual Model. In *IJCAI*, 884–890.
- Du, Y.; Chen, Z.; Su, Y.; Jia, C.; and Jiang, Y. 2024a. Instruction-Guided Scene Text Recognition. *CoRR*, abs/2401.17851.
- Du, Y.; Chen, Z.; Xie, H.; Jia, C.; and Jiang, Y.-G. 2024b. SVTRv2: CTC Beats Encoder-Decoder Models in Scene Text Recognition. *CoRR*, abs/2411.15858.
- Fang, S.; Mao, Z.; Xie, H.; Wang, Y.; Yan, C.; and Zhang, Y. 2023. ABINet++: Autonomous, Bidirectional and Iterative Language Modeling for Scene Text Spotting. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(6): 7123–7141.
- Fang, S.; Xie, H.; Wang, Y.; Mao, Z.; and Zhang, Y. 2021. Read Like Humans: Autonomous, Bidirectional and Iterative Language Modeling for Scene Text Recognition. In *CVPR*, 7098–7107.
- Graves, A.; Fernández, S.; Gomez, F.; and Schmidhuber, J. 2006. Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks. In *ICML*, 369–376.
- He, M.; Liu, Y.; Yang, Z.; Zhang, S.; Luo, C.; Gao, F.; Zheng, Q.; Wang, Y.; Zhang, X.; and Jin, L. 2018. ICDAR2018 contest on robust reading for multi-type web images. In *ICPR*, 7–12.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long Short-Term Memory. *Neural Comput.*, 9(8): 1735–1780.
- Hu, W.; Cai, X.; Hou, J.; Yi, S.; and Lin, Z. 2020. GTC: Guided training of ctc towards efficient and accurate scene text recognition. In *AAAI*, 11005–11012.
- Jiang, Q.; Wang, J.; Peng, D.; Liu, C.; and Jin, L. 2023. Revisiting Scene Text Recognition: A Data Perspective. In *ICCV*, 20486–20497.
- Jung, J.; Lee, S.; Cho, M.; and Kim, J. 2011. Touch TT: Scene text extractor using touchscreen interface. *ETRI*, 78–88.
- Karatzas, D.; Gomez-Bigorda, L.; Nicolaou, A.; Ghosh, S.; Bagdanov, A.; Iwamura, M.; Matas, J.; Neumann, L.; Chandrasekhar, V. R.; Lu, S.; Shafait, F.; Uchida, S.; and Valveny, E. 2015. ICDAR 2015 competition on Robust Reading. In *ICDAR*, 1156–1160.
- KaratzasAU, D.; ShafaitAU, F.; UchidaAU, S.; IwamuraAU, M.; i. BigordaAU, L. G.; MestreAU, S. R.; MasAU, J.; MottaAU, D. F.; AlmazànAU, J. A.; and de las Heras, L. P. 2013. ICDAR 2013 Robust Reading Competition. In *ICDAR*, 1484–1493.
- Krylov, I.; Nosov, S.; and Sovrasov, V. 2021. Open images v5 text annotation and yet another mask text spotter. In *Asian Conference on Machine Learning*, 379–389. PMLR.
- Lee, C.; and Osindero, S. 2016. Recursive recurrent nets with attention modeling for ocr in the wild. In *CVPR*, 2231–2239.
- Li, C.; Liu, W.; Guo, R.; Yin, X.; Jiang, K.; Du, Y.; Du, Y.; Zhu, L.; Lai, B.; Hu, X.; Yu, D.; and Ma, Y. 2022. PP-OCRv3: More Attempts for the Improvement of Ultra Lightweight OCR System. *CoRR*, abs/2206.03001.
- Li, H.; Wang, P.; Shen, C.; and Zhang, G. 2019. Show, attend and read: A simple and strong baseline for irregular text recognition. In *AAAI*, 8610–8617.
- Liu, Y.; Jin, L.; Zhang, S.; and Zhang, S. 2017. Detecting curve text in the wild: New dataset and new solution. *CoRR*, abs/1712.02170.
- Long, S.; Qin, S.; Panteleev, D.; Bissacco, A.; Fujii, Y.; and Raptis, M. 2022. Towards end-to-end unified scene text detection and layout analysis. In *CVPR*, 1049–1059.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Lu, N.; Yu, W.; Qi, X.; Chen, Y.; Gong, P.; Xiao, R.; and Bai, X. 2021. MASTER: Multi-aspect non-local network for scene text recognition. *Pattern Recognit.*, 117: 107980.
- Luo, C.; Jin, L.; and Sun, Z. 2019. MORAN: A Multi-Object Rectified Attention Network for Scene Text Recognition. *Pattern Recognit.*, 90: 109–118.
- Mathew, M.; Jain, M.; and Jawahar, C. 2017. Benchmarking scene text recognition in Devanagari, Telugu and Malayalam. In *ICDAR*, volume 7, 42–46.
- Mishra, A.; Karteek, A.; and Jawahar, C. V. 2012. Scene Text Recognition using Higher Order Language Priors. In *BMVC*, 1–11.

- Na, B.; Kim, Y.; and Park, S. 2022. Multi-modal Text Recognition Networks: Interactive Enhancements Between Visual and Semantic Features. In *ECCV*, 446–463.
- Nagy, R.; Dicker, A.; and Meyer-Wegener, K. 2012. NEOCR: A configurable dataset for natural image text recognition. In *Camera-Based Document Analysis and Recognition*, 150–163.
- Nayef, N.; Patel, Y.; Busta, M.; Chowdhury, P.; Karatzas, D.; Khlif, W.; Matas, J.; Pal, U.; Burie, J.; Liu, C.; et al. 2019. ICDAR2019 robust reading challenge on multi-lingual scene text detection and recognition—RRC-MLT-2019. In *ICDAR*, 1582–1587.
- Phan, T. Q.; Shivakumara, P.; Tian, S.; and Tan, C. L. 2013. Recognizing Text with Perspective Distortion in Natural Scenes. In *CVPR*, 569–576.
- Qiao, Z.; Zhou, Y.; Wei, J.; Wang, W.; Zhang, Y.; Jiang, N.; Wang, H.; and Wang, W. 2021. PIMNet: a parallel, iterative and mimicking network for scene text recognition. In *ACM MM*, 2046–2055.
- Qiao, Z.; Zhou, Y.; Yang, D.; Zhou, Y.; and Wang, W. 2020. SEED: Semantics Enhanced Encoder-Decoder Framework for Scene Text Recognition. In *CVPR*, 13525–13534.
- Sheng, F.; Chen, Z.; and Xu, B. 2019. NRTR: A No-Recurrence Sequence-to-Sequence Model for Scene Text Recognition. In *ICDAR*, 781–786.
- Shi, B.; Bai, X.; and Yao, C. 2017. An End-to-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(11): 2298–2304.
- Shi, B.; Yang, M.; Wang, X.; Lyu, P.; Yao, C.; and Bai, X. 2019. ASTER: An Attentional Scene Text Recognizer with Flexible Rectification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(9): 2035–2048.
- Shi, B.; Yao, C.; Liao, M.; Yang, M.; Xu, P.; Cui, L.; Belongie, S.; Lu, S.; and Bai, X. 2017. ICDAR2017 competition on reading chinese text in the wild (rctw-17). In *ICDAR*, 1429–1434.
- Singh, A.; Pang, G.; Toh, M.; Huang, J.; Galuba, W.; and Hassner, T. 2021. TextOCR: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text. In *CVPR*, 8802–8812.
- Sun, Y.; Ni, Z.; Chng, C.; Liu, Y.; Luo, C.; Ng, C.; Han, J.; Ding, E.; Liu, J.; Karatzas, D.; et al. 2019. ICDAR 2019 competition on large-scale street view text with partial labeling-RRC-LSVT. In *ICDAR*, 1557–1562.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NIPS*, 5998–6008.
- Veit, A.; Matera, T.; Neumann, L.; Matas, J.; and Belongie, S. 2016. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *CoRR*, abs/1601.07140.
- I. Loshchilov; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. In *ICLR*.
- Wang, K.; Babenko, B.; and Belongie, S. 2011. End-to-end scene text recognition. In *ICCV*, 1457–1464.
- Wang, T.; Zhu, Y.; Jin, L.; Luo, C.; Chen, X.; Wu, Y.; Wang, Q.; and Cai, M. 2020. Decoupled Attention Network for Text Recognition. In *AAAI*, 12216–12224.
- Wang, Y.; Xie, H.; Fang, S.; Wang, J.; Zhu, S.; and Zhang, Y. 2021. From Two to One: A New Scene Text Recognizer With Visual Language Modeling Network. In *ICCV*, 14194–14203.
- Wei, J.; Zhan, H.; Lu, Y.; Tu, X.; Yin, B.; Liu, C.; and Pal, U. 2024. Image as a Language: Revisiting Scene Text Recognition via Balanced, Unified and Synchronized Vision-Language Reasoning Network. In *AAAI*, 5885–5893.
- Xu, J.; Wang, Y.; Xie, H.; and Zhang, Y. 2024. OTE: Exploring Accurate Scene Text Recognition Using One Token. In *CVPR*, 28327–28336.
- Yang, J.; Li, C.; Dai, X.; and Gao, J. 2022. Focal Modulation Networks. In *NeurIPS*.
- Yang, M.; Yang, B.; Liao, M.; Zhu, Y.; and Bai, X. 2024. Class-Aware Mask-guided feature refinement for scene text recognition. *Pattern Recognition*, 149: 110244.
- Yu, D.; Li, X.; Zhang, C.; Liu, T.; Han, J.; Liu, J.; and Ding, E. 2020. Towards accurate scene text recognition with semantic reasoning networks. In *CVPR*, 12113–12122.
- Yu, H.; Wang, X.; Li, B.; and Xue, X. 2023. Chinese Text Recognition with A Pre-Trained CLIP-Like Model Through Image-IDS Aligning. In *ICCV*, 11909–11918.
- Yue, X.; Kuang, Z.; Lin, C.; Sun, H.; and Zhang, W. 2020. RobustScanner: Dynamically enhancing positional clues for robust text recognition. In *ECCV*, 135–151.
- Zhang, B.; Xie, H.; Wang, Y.; Xu, J.; and Zhang, Y. 2023. Linguistic More: Taking a Further Step toward Efficient and Accurate Scene Text Recognition. In *IJCAI*, 1704–1712.
- Zhang, R.; Zhou, Y.; Jiang, Q.; Song, Q.; Li, N.; Zhou, K.; Wang, L.; Wang, D.; Liao, M.; Yang, M.; et al. 2019. ICDAR 2019 robust reading challenge on reading chinese text on signboard. In *ICDAR*, 1577–1581.
- Zhang, S.; Yang, C.; Zhu, X.; Zhou, H.; Wang, H.; and Yin, X. 2024a. Inverse-Like Antagonistic Scene Text Spotting via Reading-Order Estimation and Dynamic Sampling. *IEEE Trans. Image Process.*, 33: 825–839.
- Zhang, Y.; Gueguen, L.; Zharkov, I.; Zhang, P.; Seifert, K.; and Kadlec, B. 2017. Uber-text: A large-scale dataset for optical character recognition from street-level imagery. In *Scene Understanding Workshop-CVPR*, 5.
- Zhang, Z.; Lu, N.; Liao, M.; Huang, Y.; Li, C.; Wang, M.; and Peng, W. 2024b. Self-Distillation Regularized Connectionist Temporal Classification Loss for Text Recognition: A Simple Yet Effective Approach. In *AAAI*, 7441–7449.
- Zhao, S.; Du, Y.; Chen, Z.; and Jiang, Y.-G. 2024. Decoder Pre-Training with only Text for Scene Text Recognition. In *ACM MM*, 5191–5200.
- Zheng, T.; Chen, Z.; Fang, S.; Xie, H.; and Jiang, Y. 2024. CDistNet: Perceiving multi-domain character distance for robust text recognition. *Int. J. Comput. Vis.*, 132(2): 300–318.