

Self-supervised Implicit Glyph Attention for Text Recognition

Tongkun Guan¹, Chaochen Gu^{2*}, Jingzheng Tu², Xue Yang¹, Qi Feng², Yudi Zhao², Wei Shen^{1*}

¹ MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

² Department of Automation, Shanghai Jiao Tong University

{gtk0615, jacygu, tujingzheng, yangxue-2019-sjtu, fengqi, yudizhao, wei.shen}@sjtu.edu.cn

Abstract

The attention mechanism has become the de facto module in scene text recognition (STR) methods, due to its capability of extracting character-level representations. These methods can be summarized into implicit attention based and supervised attention based, depended on how the attention is computed, i.e., implicit attention and supervised attention are learned from sequence-level text annotations and or character-level bounding box annotations, respectively. Implicit attention, as it may extract coarse or even incorrect spatial regions as character attention, is prone to suffering from an alignment-drifted issue. Supervised attention can alleviate the above issue, but it is character category-specific, which requires extra laborious character-level bounding box annotations and would be memory-intensive when handling languages with larger character categories. To address the aforementioned issues, we propose a novel attention mechanism for STR, self-supervised implicit glyph attention (SIGA). SIGA delineates the glyph structures of text images by jointly self-supervised text segmentation and implicit attention alignment, which serve as the supervision to improve attention correctness without extra character-level annotations. Experimental results demonstrate that SIGA performs consistently and significantly better than previous attention-based STR methods, in terms of both attention correctness and final recognition performance on publicly available context benchmarks and our contributed contextless benchmarks.

1. Introduction

Scene text recognition (STR) aims to recognize texts from natural images, which has wide applications in handwriting recognition [36, 47, 53], industrial print recognition [12, 16, 32], and visual understanding [8, 23, 31]. Recently, attention-based models with encoder-decoder architectures are typically developed to address this task by attending to important regions of text images to extract




Figure 1. Three difference supervised manners for STR.

character-level representations. These methods can be summarized into implicit attention methods (a) and supervised attention methods (b) as shown in Figure 1, according to the annotation type used for supervising the attention.

Specifically, implicit attention is learned from sequence-level text annotations by computing attention scores across all locations over a 1D or 2D space. For example, the 1D sequential attention weights [3, 39] are generated at different decoding steps to extract important items of the encoded sequence. The 2D attention weights [15, 50] are generated by executing a cross-attention operation with the embedded time-dependent sequences and visual features on all spatial locations. However, implicit attention methods, which only extract coarse or even unaligned spatial regions as character attention, may encounter alignment-drifted attention. In contrast, supervised attention is learned from extra character-level bounding box annotations by generating character segmentation maps. Although these supervised attention methods [20, 28, 30, 42] can alleviate the above issue, they rely on labour-intensive character-level bounding box annotations, and their attention maps with respect to character categories might be memory-intensive when the number of character categories is large.

To address the aforementioned issues, we propose a novel attention-based method for STR (Figure 1 (c)), self-supervised implicit glyph attention (SIGA). As briefly shown in Figure 2, SIGA delineates the glyph structures of text images by jointly self-supervised text segmentation and implicit attention alignment, which serve as the supervision for learning attention maps during training to improve attention correctness. Specifically, the glyph structures are generated by modulating the learned text foreground representations with sequence-aligned attention

*Corresponding author.

vectors. The text foreground representations are distilled from self-supervised segmentation results according to the internal structures of images [18]; The sequence-aligned attention vectors are obtained by applying an orthogonal constraint to the 1D implicit attention vectors [3]. They then serve as the position information of each character in a text image to modulate the text foreground representations to generate glyph pseudo-labels online.

By introducing glyph pseudo-labels as the supervision of attention maps, the learned glyph attention encourages the text recognition network to focus on the structural regions of glyphs to improve attention correctness. Different from supervised attention methods, the glyph attention maps bring no additional cost to enable character and decoding order consistency when handling languages with larger character categories.

For recognizing texts with linguistic context, SIGA achieves state-of-the-art results on seven publicly available context benchmarks. We also encapsulate our glyph attention module as a plug-in component to other attention-based methods, achieving average performance gains of 5.68% and 1.34% on SRN [50] and ABINet [15], respectively.

It is worth mentioning that SIGA shows its prominent superiority in recognizing contextless texts widely used in industrial scenarios (*e.g.*, workpiece serial numbers [16] and identification codes [32]). Specifically, we contribute two large-scale contextless benchmarks (real-world MPSC and synthetic ArbitText) with random character sequences that differ from legal words. Experiments demonstrate that SIGA improves the accuracy of contextless text recognition by a large margin, which is 7.0% and 10.3% higher than MGP-STR [44] on MPSC and ArbitText, respectively. In summary, the main contributions are as follows:

- We propose a novel attention mechanism for scene text recognition, SIGA, which is able to delineate glyph structures of text images by jointly self-supervised text segmentation and implicit attention alignment to improve attention correctness without character-level bounding box annotations.
- Extensive experiments demonstrate that the proposed glyph attention is essential for improving the performance of vision models. Our method achieves the state-of-the-art performance on publicly available context benchmarks and our contributed large-scale contextless benchmarks (MPSC and ArbitText).

2. Related Work

Recently, some top-down approaches have been developed to recognize entire images instead of directly recognizing character segments like traditional bottom-up approaches [35, 43]. These methods can be roughly divided into language-free and language-aware methods.

Language-free methods These methods view STR as a

character-level classification task and mainly exploit the visual information to recognize texts. According to the annotation types used for supervising the attention, implicit attention methods are developed for STR supervised by sequence-level text annotations, while supervised attention methods require additional character-level bounding box annotations.

Specifically, some 1D implicit attention methods [4, 11, 29, 38, 39] execute sequence attention modeling over a 1D space. Input text images are first encoded into 1D sequential features. Then, they employ a bidirectional decoder to extract attentive features of the encoded sequence by outputting the corresponding attention weights for prediction. Besides, some 2D implicit attention methods [15, 26, 27, 45, 50] develop various 2D attention mechanisms by attending to spatial vision features of each character of an image. For example, Li *et al.* [27] combine visual features with hidden states of the decoder to focus on spatial character features at each decoding step. Fang *et al.* [15] adopt a transformer-based structure to compute attention scores across all spatial locations of visual features, thereby obtaining attention maps of corresponding characters. However, supervised by sequence-level text annotations, these implicit attention methods easily extract coarse or even unaligned spatial regions as character attention.

In contrast, under the supervision of extra character-level bounding box annotations, some supervised attention methods [20, 30, 42] employ a fully convolutional network to predict character-level segmentation results and then perform classification tasks. For example, He *et al.* [20] utilizes the segmentation probability maps to exploit spatial context for text reasoning by graph convolutional networks. However, character-level annotations of text images are expensive and laborious. Beyond the limitations of human annotations, we delineate the glyph structures of text images as the supervision of attention maps by jointly self-supervised text segmentation and implicit attention alignment.

Language-aware methods Inspired by natural language processing methods [9, 24], the visual outputs of STR methods are fed into a language model to implement recognition correction with linguistic context. For example, some works [7, 50] stack multiple layers of self-attention structures [41] for semantic reasoning tasks. Inspired by the masked language model (MLM) in BERT [24], Fang *et al.* [15] pre-train the proposed BCN to predict the masked character in text based on linguistic context, and unite visual outputs to improve performance. Although these language-aware methods leverage a language model to optimize the joint character prediction probability with visual models, which reduces prediction errors with linguistic context, they do not generalize well to arbitrary texts (*e.g.*, contextless texts with a specific workpiece coding scheme). Therefore extracting the distinctive visual features of characters is still

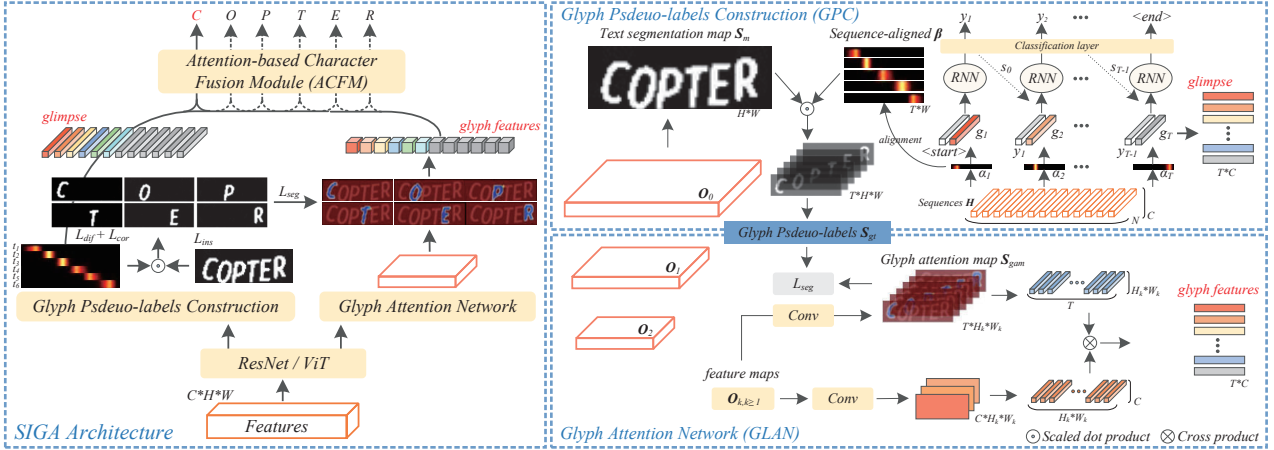


Figure 2. Overview of the proposed self-supervised implicit glyph attention network (SIGA) for text recognition.

the key to text recognition.

3. Methodology

In this section, we first review the representative attention-based method [3] that implicitly learns the 1D attention weights, and then introduce our self-supervised implicit glyph attention method.

3.1. Implicit Attention Method over 1D Space

The implicit attention method [3] consists of a transformation layer, an encoder and a decoder. First, the transformation layer employs a Thin Plate Spline (TPS), a variant of the spatial transformation network (STN) [22], to transform an input image X into a normalized image X' . Then, the encoder extracts sequential features $H \in \mathbb{R}^{C \times 1 \times N}$ from the normalized image $X' \in \mathbb{R}^{C \times H \times W}$ by a variant of ResNet [19], and splits the sequential features into a fixed-length sequence $\{h_i\}_{i=1, \dots, N}$. In the decoder, as illustrated in Figure 4 (a), the encoded sequence is fed into a recurrent module (e.g., LSTM, GRU) to generate an output vector x_t and a new state vector s_t at the decoding step t . The specific details are as follows:

$$(x_t, s_t) = \text{rnn}(s_{t-1}, (g_t, E(y_{t-1}))), \quad (1)$$

where $(g_t, E(y_{t-1}))$ denotes the combination of *glimpse* g_t and the embedding vector of the predicted character category at the previous decoding step. Especially, y_0 denotes an artificially defined “<start>” token. The *glimpse* is computed by the attention mechanism as follows:

$$\begin{cases} e_{t,i} = w^\top \tanh(Ws_{t-1} + Vh_i + b), \\ \alpha_{t,i} = \exp(e_{t,i}) / \sum_{i'=1}^N \exp(e_{t,i'}), \\ g_t = \sum_i \alpha_{t,i} h_i, i = 1, \dots, N, \end{cases} \quad (2)$$

where the w, W, V are learnable parameters. Finally, the output vector x_t predicts the character classification by a

linear layer at the current decoding step t . The decoder is executed T times (i.e., total decoding steps, $T = 26$) and outputs classification results sequentially.

3.2. Our Self-supervised Implicit Glyph Attention

In this work, we follow the implicit attention method [3] as the baseline structure, and delineate glyph structures of text images as the supervision of our attention network by proposing a novel online glyph pseudo-label construction module. The learned glyph attention encourages the text recognition network to focus on the structural regions of glyphs to improve attention correctness.

3.2.1 Glyph Pseudo-label Construction (GPC)

For a normalized image, given its text mask and the horizontal position information of each character, we can easily obtain the glyph structures of these characters by computing the dot product between them, instead of labour-intensive pixel-level annotations. Towards the goal, we construct glyph pseudo-labels online by jointly *self-supervised text segmentation* and *implicit attention alignment*. The sequence-aligned attentions serve as the position information of characters in a text image to modulate the learned text foreground representations to generate significant glyph pseudo-labels.

1) Self-supervised Text Segmentation. In the subsection, we want to learn text foreground representations with morphological structures of glyphs, by a semantic segmentation network that assigns every pixel a foreground or background label on the unlabeled text images. It’s observed that the underlying morphological representations of glyphs are not affected by slight structural changes (e.g., thicker or thinner), which reduces the reliance on pixel-level high-precision segmentation with expensive computation and annotation costs. Inspired by the prior knowledge, we begin with a clustering task based on the internal structures of text



Figure 3. Some segmentation examples. (a) denotes the original images, (b) means the results of K -means, and (c) is our text segmentation maps. In the last group of rows, the well-learned self-supervised text segmentation module can capture morphological structures of glyphs in challenging images.

images for obtaining pseudo-labels S_{pl} about text masks. For simplicity, we focus our study on K -means to implement the clustering task, but other clustering approaches with predefined categories can be used. In the experiment, K is set to 2, including the foreground and background categories. Surprisingly, the morphological structures of glyphs are clustered well in most text images.

Then, the text foreground representations are distilled from self-supervised segmentation results produced by our designed text segmentation network. Specifically, we define the output of *Conv 0*, *Block 0*, and *Block 1* from ResNet as P_0 , P_1 , and P_2 , and a top-down pyramid architecture is employed as follows:

$$\begin{cases} O_2 = \varphi(P_2), \\ O_1 = \varphi([\mathcal{T}(O_2, s_1), P_1]), \\ O_0 = \varphi([\mathcal{T}(O_1, s_0), P_0]), \end{cases} \quad (3)$$

where $\varphi(\cdot)$ denotes two convolutional layers with Batch-Norm and ReLU activation function, $\mathcal{T}(\cdot)$ refers to a single $2 \times$ upsampling for O_k with resolution s_k (i.e., $H_k \times W_k$), and $[\cdot]$ represents the concatenation operation along the channel axis. O_0 is exploited to produce the text segmentation mask S_m by a binary classification convolutional layer.

Finally, we employ a binary cross-entropy loss \mathcal{L}_{ins} between the text segmentation mask S_m and pseudo-labels S_{pl} to optimize the text segmentation network. Consequently, the optimized segmentation network perceives the text foreground representations with morphological structures of glyphs in challenging text images, which may be difficult to be classified by an unsupervised clustering method K -means. Some visualization examples are shown in Figure 3.

2) Implicit Attention Alignment. In the decoding unit (Eq. 2), the implicit attention weights $\alpha = \{\alpha_t\}_{t=1, \dots, T}$ focus on the important items of the encoded sequence to capture character dependencies. Inspiringly, we transform the attention weights as the position information of their corresponding characters. However, the time information of the decoder is drowned with the other introductions at the latter decoding steps, which easily leads to alignment drift as shown in Figure 4 (b), i.e., the learnable attention weights




Figure 4. Illustration of the representative attention-based decoder and some sequence-unaligned attention examples. The red dashed boxes in (b) indicate that the attention weight struggles to align text sequence at the current decoding step.

struggle to align the text sequence [45].

To address the issue, we apply an orthogonal constraint to the implicit attention weights to obtain sequence-aligned attention vectors. Specifically, we take these learnable attention weights as vectors and perform an alignment operation by ensuring that they are orthogonal to each other and that each processed vector is aligned with the corresponding character of the text segmentation mask S_m . Assuming that L denotes the character number of a text image, we first calculate the correlation coefficient S_{cor} between L attention vectors, and then extract the character saliency map S_{sal} by the attention vectors. The details are as follows:

$$\begin{cases} S_{cor} = \sum_{1 \leq t < t' \leq L} \alpha_t^T \alpha_{t'}, \\ \beta_t = \xi(\alpha_t), \\ S_{sal} = \sum_{t=1}^L (\sigma(\beta_t) \cdot S_m), \end{cases} \quad (4)$$

where ξ represents the one-dimensional linear interpolation ($\xi : \alpha_t \in \mathbb{R}^N \rightarrow \beta_t \in \mathbb{R}^W$). $\sigma(\cdot)$ refers to the nonlinear activation function, which maps each element in the vector to $[0, 1]$:

$$\sigma(x) = 1 / (1 + \exp(-\mu(x - \lambda))), \quad (5)$$

where μ, λ represent scaling and offset transitions, set to 70 and 0.1 in the experiment, respectively. Then the alignment drift problem can be alleviated by minimizing the correlation coefficient S_{cor} and the difference S_{dif} between S_m and S_{sal} by the following loss function:

$$\begin{cases} \mathcal{L}_{cor} = S_{cor}, \\ \mathcal{L}_{dif} = \frac{1}{n} \sum_{i=1}^n -(\rho_i \log(\rho_i^*) + (1 - \rho_i) \log(1 - \rho_i^*)), \\ \mathcal{L}_{seq} = \mathcal{L}_{cor} + \mathcal{L}_{dif}, \end{cases} \quad (6)$$

where n denotes the number of pixels in the text segmentation map S_m , ρ_i and ρ_i^* are the confidence score of pixel i in S_m and S_{sal} , respectively.

Finally, by optimizing the proposed constraint function during training, the attention weights are successfully aligned with the encoded sequence and contribute accurate positional information for glyph pseudo-label construction.

3) Glyph Pseudo-label Construction. By calculating




Figure 5. Self-constructed glyph pseudo-labels online.

the dot product between the aligned attention weights $\beta = \{\beta_t\}_{t=1, \dots, T}$ and text segmentation mask \mathcal{S}_m , we obtain the glyph pseudo-labels. Specifically, assuming that the glyph pseudo-label is \mathcal{S}_{gt} , we construct it from the concatenation operation as follows:

$$\mathcal{S}_{gt} = [1 - \mathcal{S}_m, \mathbb{1}_{[\beta_1 \geq \delta]} \cdot \mathcal{S}_m, \dots, \mathbb{1}_{[\beta_T \geq \delta]} \cdot \mathcal{S}_m], \quad (7)$$

where $[\cdot]$ represents the concatenation operation along the channel axis. δ denotes the confidence threshold, which is set to 0.05 in the experiment. Some visualization examples about the self-constructed \mathcal{S}_{gt} are shown in Figure 5.

Note that the proposed glyph pseudo-label construction module will be removed in the test stage.

3.2.2 Glyph Attention Network (GLAN)

The existing supervised attention methods for STR have the following limitations: 1) For languages with larger character categories, these methods might be memory-intensive and run slower due to their category-dependence character segmentation maps. 2) It is not easy to obtain the character order of text directly from the character segmentation maps predicted by CNNs. An extra order segmentation branch is usually introduced to ensure channel and decoding order consistency, which brings time and computational complexity. 3) Training the segmentation network requires laborious and difficult character-level bounding box annotations.

Benefiting from self-constructed glyph pseudo-labels, our glyph attention network does not have these limitations since the order is ensured and characters are well aligned. The glyph attention network generates glyph attention maps with fixed-length and category-independent channels, whose channel-specific map corresponds to the order-specific glyph attention. Specifically, followed by several convolutional layers, the features \mathcal{O}_k in Eq. 3 are utilized to predict a glyph attention map \mathcal{S}_{gam} with a channel number of N_s . N_s is set to $1 + M$ and not equal to the character categories, which represents the sum of the background category and the set maximum character length on text images ($M = 26$). For example, to recognize GB2312 with 6763 categories in Chinese, if the same convolution layer is employed and feature channels are 256, the parameter size is 1.7M (256×6763) for supervised attention methods while 6.9K (256×27) for our method.

And then, supervised by the constructed glyph pseudo-labels \mathcal{S}_{gt} , we use the joint loss function of multi-class Dice loss [33] and cross-entropy loss to boost the segmentation performance of the glyph attention network. The specific

Table 1. The parameter setting table of SIGA.

Name	Value	Name	Value
Decoding step T	26	Sequence length N	32
Feature channel C	256	Max character length M	26
Image size W, H	128, 32	Constant λ, μ	0.1, 70
Confidence threshold δ	0.05	Constant k	2

details are as follows:

$$\begin{cases} \mathcal{L}_{dice} = \frac{1}{L} \sum_{j=2}^{L+1} \left(1 - \frac{2 \sum_{i=1}^n (\omega_{j,i} \omega_{j,i}^*)}{\sum_{i=1}^n (\omega_{j,i}) + \sum_{i=1}^n (\omega_{j,i}^*)} \right), \\ \mathcal{L}_{cel} = \frac{-1}{n} \sum_{i=1}^n (\rho_i \log(\sum_{j=2}^{M+1} \omega_{j,i}^*) + (1 - \rho_i) \log(1 - \sum_{j=2}^{M+1} \omega_{j,i}^*)), \\ \mathcal{L}_{seg} = \mathcal{L}_{dice} + \mathcal{L}_{cel}, \end{cases} \quad (8)$$

where $\omega_{j,i}$ and $\omega_{j,i}^*$ are the confidence scores of the i -th pixel p_i of the j -index map in the pseudo-label \mathcal{S}_{gt} and glyph attention map \mathcal{S}_{gam} , respectively. ρ_i is the confidence score of the p_i in \mathcal{S}_m . L denotes the character number of a text image.

Finally, the learned glyph attention encourages the recognition branch to focus on the structural regions of glyphs to extract glyph features for STR, which contain more robust and discerning character representations. Specifically, the encoded text features $\mathcal{O}_k \in \mathbb{R}^{W_k \times H_k \times C}$ are first fed into two convolutional layers with BatchNorm and ReLU activation functions, and then multiplied with glyph attention maps $\mathcal{S}_{gam} \in \mathbb{R}^{W_k \times H_k \times M}$ (remove background) to obtain glyph features $\mathcal{I}_k \in \mathbb{R}^{M \times C}$.

3.2.3 Attention-based Character Fusion Module

As discussed above, the visually aligned *glimpse* \mathbf{g}_t and glyph features $\mathcal{I}_{k,t}$ denote two different character feature representations at the decoding step t . Considering that their contributions to STR should be different among various text images, inspired by the gate unit [1], we dynamically fuse the sequence $\mathcal{I}_{k,t}$ and *glimpse* \mathbf{g}_t to enrich the semantic information for character recognition. Finally, we embed the final sequence into a decoder [3] to output the current decoded classification result.

4. Experiments

4.1. Datasets

Our model is trained on two large-scale synthetic datasets (*i.e.* SynthText [17] and MJSynth [21]) for a fair comparison. Nine STR datasets are used to evaluate the performance of our method, including seven publicly available context benchmarks [3] (*i.e.*, IIIT5K-Words, ICDAR2003, ICDAR2013, Street View Text, ICDAR2015, SVT Perspective, and CUTE80) and two contextless benchmarks (MPSC and ArbitText). The differences between context and contextless benchmarks are shown in Figure 6.

Table 2. Comparison results of language-free STR methods. † represents the visual model performance for a fair comparison. * combines 21 mixing blocks (10 local blocks and 11 global blocks) with local and global modeling capabilities for extracting features. “Trns” refers to several transformer units [41] consisting of a MHSA and a FFN. “SATRN” is tailored for a transformer-based text feature extractor [26]. These symbols follow the same convention within the scope of this paper. The best results are shown in bold font. Underline values represent the second-best results.

Methods	Backbone	Structure	Size	Venue	IIT 3000	SVT 647	IC03 860 867	IC13 857 1015	IC15 1811 2077	SP 645	CT 288		
CA-FCN [28]	VGG16	CNN	64×256	AAAI2019	91.9	86.4	-	-	91.5	-	-	79.9	
DAN [45]	ResNet45		32×128	AAAI2020	94.3	89.2	-	95.0	-	93.9	-	80.0	
TextScanner [42]	ResNet50		64×256	AAAI2020	93.9	90.1	-	-	92.9	79.4	-	84.3	
SRN† [50]	ResNet50-FPN		64×256	CVPR2020	92.3	88.1	-	-	93.2	77.5	-	84.7	
PlugNet [34]	ResNet37		32×100	ECCV2020	94.4	92.3	95.7	-	95.0	-	82.2	84.3	
PIMNet [37]	ResNet50-FPN		64×256	ACM MM2021	95.2	91.2	-	-	95.2	93.4	83.5	81.0	
TRBA [3]	ResNet31		32×100	CVPR2021	92.1	88.9	94.8	<u>95.1</u>	93.9	93.1	78.3	74.7	
PREN2D [49]	EfficientNet-B3		64×256	CVPR2021	<u>95.6</u>	94.0	<u>95.8</u>	-	96.4	-	83.0	-	
Text is Text [6]	ResNet31		48×160	ICCV2021	92.3	89.9	-	-	93.3	-	-	76.9	
S-GTR† [20]	ResNet50Dilated-PPM		64×256	AAAI2022	94.0	91.2	-	-	94.8	-	82.8	-	
LevOCR† [10]	ResNet45		32×128	ECCV2022	95.2	90.6	-	-	95.1	-	<u>84.0</u>	-	
SGBANet [54]	ResNet45-FPN		64×256	ECCV2022	95.4	89.1	-	-	-	<u>95.1</u>	-	78.4	
SIGAR (ours)	ResNet45		32×128	-	-	95.9	<u>92.7</u>	96.5	95.9	97.0	95.6	85.1	81.7
SVTR [14]	SVTR-L		Transformer*	48×160	IJCAI2022	<u>96.3</u>	<u>91.7</u>	-	-	97.2	-	<u>86.6</u>	-
SIGAS (ours)	SVTR-L			48×160	-	-	96.9	93.7	-	-	<u>97.0</u>	-	87.6
ViTSTR [2]	ViT-B		Transformer	224×224	ICDAR2021	88.4	87.7	<u>94.7</u>	94.3	93.2	92.4	78.5	72.6
ABINet† [15]	ResNet45-Trns	32×128		CVPR2021	94.7	91.7	94.3	<u>94.7</u>	95.0	93.6	82.7	<u>83.0</u>	
ABINet+ConCLR† [52]	ResNet45-Trns	32×128		AAAI2022	95.7	92.1	-	-	95.9	93.6	84.4	-	
LevOCR† [10]	ViT	32×128		ECCV2022	93.6	89.2	-	-	94.9	-	82.4	-	
CornerTransformer [48]	SATRN	32×128		ECCV2022	95.9	94.6	-	-	-	<u>96.4</u>	-	86.3	
MGP-STR [44]	ViT-B	32×128		ECCV2022	<u>96.4</u>	<u>94.7</u>	-	-	<u>97.3</u>	-	87.2	-	
SIGAT (ours)	ViT-B	32×128		-	-	96.6	95.1	96.9	97.0	97.8	96.8	<u>86.6</u>	<u>83.0</u>
													90.5

Table 3. Comparison results of language-aware STR methods. “V” and “VL” types refer to language-free model and language-aware model, respectively. The best results are shown in bold font. Underline values represent the second-best results.

Methods	Types	Backbone	Structure	Size	Venue	IIT 3000	SVT 647	IC13 857 1015	IC15 1811 2077	SP 645	CT 288	
SRN [50]	VL	ResNet50-FPN	ResNet	64×256	CVPR2020	94.8	91.5	-	95.5	82.7	-	
Bhunia et al. [7]	VL	ResNet50-FPN		32×100	ICCV2021	95.2	92.2	-	95.5	-	84.0	
VisionLAN [46]	VL	ResNet45		64×256	ICCV2021	95.8	91.7	-	95.7	83.7	-	
S-GTR [20]	VL	ResNet50Dilated-PPM		64×256	AAAI2022	95.8	94.1	96.8	-	84.6	-	
LevOCR [10]	VL	ResNet45		32×128	ECCV2022	96.6	<u>92.9</u>	96.9	-	86.4	-	
SIGAR (ours)	V	ResNet45		32×128	-	-	<u>95.9</u>	<u>92.7</u>	97.0	<u>95.6</u>	<u>85.1</u>	<u>81.7</u>
ABINet [15]	VL	ResNet45-Trns	Transformer	32×128	CVPR2021	96.2	93.5	<u>97.4</u>	-	86.0	-	
ABINet+ConCLR [52]	VL	ResNet45-Trns		32×128	AAAI2022	96.5	94.3	-	97.7	85.4	-	
PARSeq [5]	VL	DeiT		32×128	ECCV2022	97.0	<u>93.6</u>	97.0	96.2	<u>86.5</u>	<u>82.9</u>	
LevOCR [10]	VL	ViT		32×128	ECCV2022	95.6	91.8	96.2	-	85.8	-	
SIGAT (ours)	V	ViT-B		32×128	-	-	<u>96.6</u>	95.1	97.8	<u>96.8</u>	86.6	83.0
											90.5	93.1

MPSC: We cropped 15003 real-world text instances from industrial images marked workpiece information [16], which is larger than the sum of seven context benchmarks. These texts are randomly collected from massive internet images and not from the same batch of products, containing various workpieces with irregular character combinations (e.g., “YS6Q-6615-AD”, “TBJU8549728”, and “RS550SH-4941”) for marking workpiece information.

ArbitText: We also synthesize a contextless ArbitText with 1M images, and every sample is generated by a random combination of English letters and Arabic numerals.

4.2. Implementation Details

The parameter details are shown in Table 1. According to the backbone types of the existing text recognition models, we construct three typical SIGA architectures, i.e., SIGAR, SIGAS and SIGAT, for a fair comparison. For SIGAR with a ResNet45 [19] as the backbone, we adopt the Adam optimizer [25] and the one-cycle learning rate scheduler [40] with a maximum learning rate of 0.0005 to train our model. We employ the same augmentation strategy

from ABINet [15] and set the batch size to 512 and the training epoch to 6. For SIGAS, we select SVTR-L [14] with local and global modeling capabilities as the backbone, and the training parameters are the same as SIGAR. For SIGAT using ViT [13] as the backbone, we utilize the same setting including optimizer, learning rate scheduler, and batch size from MGP-STR [44]. To adapt the transformer structure to our method, we select the output of 2, 4, and 6 layers of ViT as P_0 , P_1 , and P_2 to execute the GPC module.

4.3. Comparisons on context benchmarks

Language-free model. The language-free methods mainly exploit visual information to recognize texts. As shown in Table 2, we compare with the previous state-of-the-art language-free methods according to backbone types to fairly evaluate the effectiveness of our method on standard context benchmarks.

For the CNN-based methods, SIGAR achieves state-of-the-art performance on seven context benchmarks. Specifically, compared to supervised attention methods (CA-FCN [28] and TextScanner [42]), SIGAR doesn’t need ex-




Figure 6. Comparison of different text benchmarks of STR.

tra character-level annotations and brings significant performance gains (2.0% ~ 11.8%) on these benchmarks. Compared to implicit attention methods, $SIGA_R$ has better performances and outperforms the second-best results on IIT, IC03-860, IC03-867, IC13-857, IC13-1015, and IC15-1811 benchmarks by 0.3%, 0.7%, 0.8%, 0.6%, 0.5%, and 1.1%, respectively. $SIGA_R$ also achieves competitive performances (underline values) on SVT, IC15-2077, and SP benchmarks.

We also deploy the backbone of SVTR [14] to implement STR, and consequently, $SIGA_S$ gets higher accuracy on four of its reported six standard benchmarks, with an average accuracy improvement of 0.63%.

For the Transformer-based methods, $SIGA_T$ shows its prominent superiority and achieves state-of-the-art results on IIT, SVT, IC03-860, IC03-867, IC13-857, IC13-1015, and CT benchmarks. Besides, we also obtain competitive results on IC15-1811 and IC15-2077 benchmarks. These results demonstrate the effectiveness of our method on context benchmarks, as more discerning visual features are successfully extracted by introducing glyph attention.

Language-aware model. The semantic reasoning task of language models corrects visual outputs to reduce prediction errors with linguistic context (e.g., correcting “university” to “university”), which improves the overall recognition accuracy on context benchmarks. As shown in Table 3, when further compared with these language-aware methods, $SIGA_R$ achieves competitive results on the most standard benchmarks, and $SIGA_T$ gets the best accuracy on six of the eight benchmarks despite not using the semantic reasoning task on benchmarks with linguistic context. Specifically, $SIGA_T$ has better performances on SVT, IC13-857, SP and CT benchmarks by 1.5%, 0.4%, 1.2%, and 0.9%, respectively, which implies that a visual model could still perform well on context benchmarks.

4.4. Comparisons on contextless benchmarks

Contextless texts consist of random character sequences, which are widely used in industrial scenarios [16, 32]. Different from context benchmarks, they contain less semantic information. Thus language-aware methods that build implicit language representations with linguistic context are unsuitable for these contextless texts. Exploiting the visual features of text images is crucial for improving recognition accuracy on contextless benchmarks. As shown in Table 4, we conduct comparative experiments with other language-free text recognition methods by utilizing the same training data (MJ and ST). These results are obtained by directly

Table 4. Comparison results of contextless benchmarks.

Methods	Venue	MPSC 15003	ArbitText 1000000
SAR [27]	AAAI2019	59.7	64.5
DAN [45]	AAAI2020	57.7	61.0
GA-SPIN [51]	AAAI2021	51.7	54.0
PIMNet [37]	ACM MM2021	60.8	61.1
$SIGA_R$ (ours)	-	65.6	66.0
SVTR [14]	IJCAI2022	71.4	78.1
$SIGA_S$ (ours)	-	72.7	81.0
ABINet† [15]	CVPR2021	64.4	61.8
MGP-STR [44]	ECCV2022	65.0	61.4
$SIGA_T$ (ours)	-	72.0	71.7




Figure 7. The distribution result of the maximum value of α .

loading their released checkpoints to be evaluated. Specifically, we first evaluate on our contributed real-world MPSC benchmark. Then, we also synthesize a large-scale contextless benchmark with 1M text images, ArbitText, to further evaluate the generality and effectiveness of language-free models.

Consequently, our method shows a significant superiority on these contextless benchmarks, as $SIGA_R$ is 4.8% and 4.9% higher than PIMNet [37], $SIGA_S$ is 1.3% and 2.9% higher than SVTR [14], $SIGA_T$ is 7.0% and 10.3% higher than MGP-STR [44] on the real-world MPSC and synthetic ArbitText benchmarks.

These results consistently emphasize that SIGA can generalize well to arbitrary texts (context benchmarks and contextless benchmarks), and the proposed glyph attention is essential for improving the performance of visual models.

4.5. Ablation Study

For efficiency, all ablation experiments are carried out by using $SIGA_R$.

Network Structure. Our architecture consists of the Glyph Pseudo-label Construction (GPC), Glyph Attention Network (GLAN), and Attention-based Character Fusion Module (ACFM). However, GLAN should be evaluated with GPC as a joint structure (JS) due to their interdependence, i.e., GLAN is supervised by glyph pseudo-labels generated from GPC during training. Thus we first perform the “Baseline+JS” model to evaluate the effectiveness of our structure. As depicted in Table 5, the accuracy of “Baseline+JS” model is 7.01% higher than “Baseline”. Then, the gain of adding ACFM is further improved by 0.45%, from 91.65% to 92.10% on average accuracy. Besides, we also encapsulate the JS structure as a plug-in component to SRN† and ABINet†. As shown in the second group of rows in Table

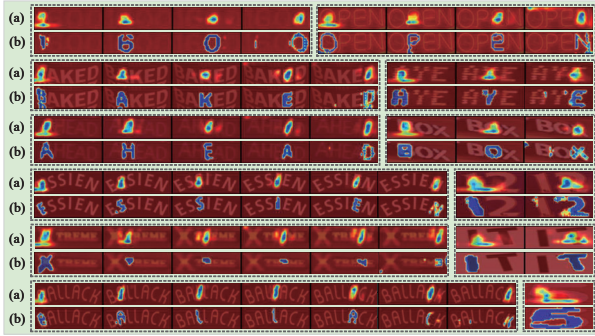


Figure 8. Extensive 2D attention visualization results generated by ABINet (a) and SIGA (b).

5, after using our JS structure, the average accuracy of these models is further improved by 5.68% and 1.34%, respectively. These results demonstrate that the glyph attention extracted by the JS structure is effective and important to facilitate character recognition.

Effectiveness of Implicit Attention Alignment. We provide the effectiveness analysis and theoretical basis about our implicit attention alignment module in Supplementary Material.

Hyper-parameters. The hyper-parameters map the attention weights α to different numerical ranges. For λ and μ of Eq. 5, they encode α to $[0, 1]$ for generating S_{sal} (Minimizing the difference between S_{sal} and S_{m} to alleviate the alignment drift problem). For δ of Eq. 7, it binarizes α to 0 or 1 for constructing S_{gt} (Obtaining the glyph pseudo-labels). To further observe the attention weight values α , we statistically analyze the distribution of the maximum value of α on the validation set, as shown in Figure 7. Thus the hyper-parameters are easily set to reasonable values based on the prior distribution. Specifically, we set the confidence threshold δ to 0.05, 0.1, and 0.15, respectively, and the average accuracy on ten standard context benchmarks is 91.19%, 90.75%, and 90.43%. For μ and λ of Eq. 5, we set the three suitable parameter pairs (μ, λ) to (100, 0.05), (70, 0.1), (40, 0.15), respectively, and the average accuracy is 91.02%, 91.19% and 90.83%.

Visualization Analysis. We have visualized the character attention maps in Figure 8, which indicates important regions contributed to character recognition. More visualization results of the SIGA method on horizontal, oriented, curved and blurred text images in Supplementary Material. Specifically, each map in (a) and (b) is generated by the representative implicit attention method ABINet [15] and our method SIGA, respectively. Unlike implicit attention mechanisms, our method perceives more fine-grained structural information of glyph. Especially, our attention can still degenerate into the same attention form as other STR models in text images with very indistinct glyph (the fifth group of rows in Figure 8). To the best of our knowledge, our method is the first to explore the glyph structures in STR.

Table 5. Ablation study of the proposed SIGA_R structure on context benchmarks. “JS” means the joint structure of GPC and GLAN, as GLAN needs GPC for glyph pseudo-label construction.

Methods	Datasets						Average
	IIIT 3000	SVT 647	IC13 857	IC15 1811	SP 645	CT 288	
Baseline	87.9	87.5	93.6	77.6	79.2	74.0	84.64
Baseline+JS	95.7	92.0	96.6	84.7	85.7	91.0	91.65
Baseline+JS+ACFM	95.9	92.7	97.0	85.1	87.1	91.7	92.10
SRN†	92.3	88.1	-	77.5	79.4	84.7	86.04
SRN†+JS	96.2	92.7	96.3	84.3	85.4	90.0	91.72
ABINet†	94.7	91.7	95.0	82.7	85.1	86.5	90.29
ABINet†+JS	95.9	92.0	96.4	84.4	85.6	91.0	91.63

Table 6. Comparison results of speed and parameter amount.

Methods	Description	Size	Param. (MB)	Time (ms)
CA-FCN [28]	Sup.	64×256	-	-
TextScanner [42]	Sup.	64×256	57	56.8
SRN† [50]	Implicit	64×256	41.4	131.5
TRBA [3]	Implicit	32×100	49.6	27.6
ABINet† [15]	Implicit	32×128	23.5	16.7
SIGA_R	Self-Sup.	32×128	40.4	53.7
CornerTransformer [48]	Implicit	32×128	85.7	294.9
LevOCR [10]	Implicit	32×128	109.0	119.0
MGP-STR [44]	Implicit	32×128	148.0	12.3
SIGA_T	Self-Sup.	32×128	113.1	56.3

Performance and Cost. We add speed and parameter comparisons in Table 6. For supervised attention methods (*i.e.*, CA-FCN and TextScanner) with more parameters due to the large input size, *e.g.*, 64×256 , SIGA_R doesn’t need extra character-level annotations and has better performance than their methods in Table 2. For implicit attention methods, SIGA_R adds an unavoidable but acceptable amount of parameters (16.9M larger than ABINet†, mainly used for a lightweight glyph attention structure) while obtaining more detailed glyph structures not being explored by other STR methods. Consequently, our method achieves the best performance in Table 2 and 4.

5. Conclusions

In this paper, we propose a novel attention-based method for STR, Self-supervised Implicit Glyph Attention (SIGA). Beyond the difficulty of character-level annotation by humans, SIGA delineates the glyph structures of text images as the supervision of attention maps by jointly self-supervised text segmentation and implicit attention alignment. The learned glyph attention then encourages the text recognition network to focus on the structural regions of glyphs to improve attention correctness. Finally, extensive experiments demonstrate that SIGA achieves the best performance on context and contextless benchmarks.

Acknowledgements This work was supported by National Key Research and Development Project of China under Grant 2019YFB1706602, NSFC 62273235, NSFC 62176159, Natural Science Foundation of Shanghai 21ZR1432200, Shanghai Municipal Science and Technology Major Project 2021SHZDZX0102 and the Fundamental Research Funds for the Central Universities.

References

- [1] John Arevalo, Tamar Solorio, Manuel Montes-y Gómez, and Fabio A González. Gated multimodal units for information fusion. *arXiv:1702.01992*, 2017. **5**
- [2] Rowel Atienza. Vision transformer for fast and efficient scene text recognition. In *ICDAR*, volume 12821, pages 319–334, 2021. **6**
- [3] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwalsuk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *ICCV*, pages 4715–4723, 2019. **1, 2, 3, 5, 6, 8**
- [4] Fan Bai, Zhazhan Cheng, Yi Niu, Shiliang Pu, and Shuigeng Zhou. Edit probability for scene text recognition. In *CVPR*, pages 1508–1516, 2018. **2**
- [5] Darwin Bautista and Rowel Atienza. Scene text recognition with permuted autoregressive sequence models. In *ECCV*, volume 13688, pages 178–196, 2022. **6**
- [6] Ayan Kumar Bhunia, Aneeshan Sain, Pinaki Nath Chowdhury, and Yi-Zhe Song. Text is text, no matter what: Unifying text recognition using knowledge distillation. In *ICCV*, pages 983–992, 2021. **6**
- [7] Ayan Kumar Bhunia, Aneeshan Sain, Amandeep Kumar, Shuvojit Ghose, Pinaki Nath Chowdhury, and Yi-Zhe Song. Joint visual semantic reasoning: Multi-stage decoder for text recognition. In *ICCV*, pages 14940–14949, 2021. **2, 6**
- [8] Ali Furkan Biten, Ruben Tito, Andres Mafla, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *ICCV*, pages 4291–4301, 2019. **1**
- [9] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020. **2**
- [10] Da Cheng, Peng Wang, and Cong Yao. Levenshtein ocr. In *ECCV*, volume 13688, pages 322–338, 2022. **6, 8**
- [11] Zhazhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. Focusing attention: Towards accurate text recognition in natural images. In *ICCV*, pages 5076–5084, 2017. **2**
- [12] Guowei Deng, Jingzheng Tu, Cailian Chen, Jianping He, and Xinyi Le. Knowledge-based scene text recognition for industrial applications. In *IEEE International Conference on Industrial Technology (ICIT)*, pages 1–6, 2022. **1**
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. **6**
- [14] Yongkun Du, Zhineng Chen, Caiyan Jia, Xiaoting Yin, Tianlun Zheng, Chenxia Li, Yuning Du, and Yugang Jiang. SVTR: scene text recognition with a single visual model. In *IJCAI*, pages 884–890, 2022. **6, 7**
- [15] 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 *CVPR*, pages 7098–7107, 2021. **1, 2, 6, 7, 8**
- [16] Tongkun Guan, Chaochen Gu, Changsheng Lu, Jingzheng Tu, Qi Feng, Kaijie Wu, and Xinpeng Guan. Industrial scene text detection with refined feature-attentive network. *IEEE TCSVT*, 32(9):6073–6085, 2022. **1, 2, 6, 7**
- [17] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *CVPR*, pages 2315–2324, 2016. **5**
- [18] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979. **2**
- [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. **3, 6**
- [20] Yue He, Chen Chen, Jing Zhang, Juhua Liu, Fengxiang He, Chaoyue Wang, and Bo Du. Visual semantics allow for textual reasoning better in scene text recognition. In *AAAI*, volume 36, pages 888–896, 2022. **1, 2, 6**
- [21] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv:1406.2227*, 2014. **5**
- [22] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. Spatial transformer networks. In *NeurIPS*, volume 28, pages 2017–2025, 2015. **3**
- [23] Zanzia Jin, Mike Zheng Shou, Fang Zhou, Satoshi Tsutsui, Jingyan Qin, and Xucheng Yin. From token to word: OCR token evolution via contrastive learning and semantic matching for text-vqa. In *ACM MM*, 2022. **1**
- [24] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019. **2**
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014. **6**
- [26] 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 *CVPRW*, pages 546–547, 2020. **2, 6**
- [27] Hui Li, Peng Wang, Chunhua Shen, and Guyu Zhang. Show, attend and read: A simple and strong baseline for irregular text recognition. In *AAAI*, pages 8610–8617, 2019. **2, 7**
- [28] Minghui Liao, Jian Zhang, Zhaoyi Wan, Fengming Xie, Jiajun Liang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Scene text recognition from two-dimensional perspective. In *AAAI*, pages 8714–8721, 2019. **1, 6, 8**
- [29] Zichuan Liu, Yixing Li, Fengbo Ren, Wang Ling Goh, and Hao Yu. Squeezedtext: A real-time scene text recognition by binary convolutional encoder-decoder network. In *AAAI*, volume 32, 2018. **2**
- [30] Pengyuan Lyu, Minghui Liao, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural

- network for spotting text with arbitrary shapes. In *ECCV*, pages 67–83, 2018. 1, 2
- [31] Mengkai Ma, Qiu-Feng Wang, Shan Huang, Shen Huang, Yannis Goulermas, and Kaizhu Huang. Residual attention-based multi-scale script identification in scene text images. *Neurocomputing*, 421:222–233, 2021. 1
- [32] Qiang Mei, Qinyou Hu, Chun Yang, Hailin Zheng, and Zhisheng Hu. Port recommendation system for alternative container port destinations using a novel neural language-based algorithm. *IEEE Access*, 8:199970–199979, 2020. 1, 2, 7
- [33] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571, 2016. 5
- [34] Yongqiang Mou, Lei Tan, Hui Yang, Jingying Chen, Leyuan Liu, Rui Yan, and Yaohong Huang. Plugnet: Degradation aware scene text recognition supervised by a pluggable super-resolution unit. In *ECCV*, pages 158–174. Springer, 2020. 6
- [35] Lukáš Neumann and Jiří Matas. Real-time scene text localization and recognition. In *CVPR*, pages 3538–3545. IEEE, 2012. 2
- [36] Dezhi Peng, Lianwen Jin, Yuliang Liu, Canjie Luo, and Songxuan Lai. Pagenet: Towards end-to-end weakly supervised page-level handwritten chinese text recognition. *IJCV*, 130(11):2623–2645, 2022. 1
- [37] Zhi Qiao, Yu Zhou, Jin Wei, Wei Wang, Yuan Zhang, Ning Jiang, Hongbin Wang, and Weiping Wang. Pimnet: a parallel, iterative and mimicking network for scene text recognition. In *ACM MM*, pages 2046–2055, 2021. 6, 7
- [38] Haibo Qin, Chun Yang, Xiaobin Zhu, and Xucheng Yin. Dynamic receptive field adaptation for attention-based text recognition. In *ICDAR*, volume 12822, pages 225–239, 2021. 2
- [39] Baoguang Shi, Mingkun Yang, Xinggong Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Aster: An attentional scene text recognizer with flexible rectification. *IEEE TPAMI*, 41(9):2035–2048, 2018. 1, 2
- [40] 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, page 1100612. International Society for Optics and Photonics, 2019. 6
- [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017. 2, 6
- [42] Zhaoyi Wan, Minghang He, Haoran Chen, Xiang Bai, and Cong Yao. Textscanner: Reading characters in order for robust scene text recognition. In *AAAI*, pages 12120–12127, 2020. 1, 2, 6, 8
- [43] Kai Wang and Serge Belongie. Word spotting in the wild. In *ECCV*, pages 591–604. Springer, 2010. 2
- [44] Peng Wang, Cheng Da, and Cong Yao. Multi-granularity prediction for scene text recognition. In *ECCV*, volume 13688, pages 339–355, 2022. 2, 6, 7, 8
- [45] Tianwei Wang, Yuanzhi Zhu, Lianwen Jin, Canjie Luo, Xiaoxue Chen, Yaqiang Wu, Qianying Wang, and Mingxiang Cai. Decoupled attention network for text recognition. In *AAAI*, pages 12216–12224, 2020. 2, 4, 6, 7
- [46] Yuxin Wang, Hongtao Xie, Shancheng Fang, Jing Wang, Shenggao Zhu, and Yongdong Zhang. From two to one: A new scene text recognizer with visual language modeling network. In *ICCV*, pages 14194–14203, 2021. 6
- [47] Jinwen Wu, Fei Yin, Yanming Zhang, Xuyao Zhang, and Chenglin Liu. Graph-to-graph: Towards accurate and interpretable online handwritten mathematical expression recognition. In *AAAI*, pages 2925–2933, 2021. 1
- [48] Xudong Xie, Ling Fu, Zhifei Zhang, Zhaowen Wang, and Xiang Bai. Toward understanding wordart: Corner-guided transformer for scene text recognition. In *ECCV*, volume 13688, pages 303–321, 2022. 6, 8
- [49] Ruijie Yan, Liangrui Peng, Shanyu Xiao, and Gang Yao. Primitive representation learning for scene text recognition. In *CVPR*, pages 284–293, 2021. 6
- [50] Deli Yu, Xuan Li, Chengquan Zhang, Tao Liu, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text recognition with semantic reasoning networks. In *CVPR*, pages 12113–12122, 2020. 1, 2, 6, 8
- [51] Chengwei Zhang, Yunlu Xu, Zhanzhan Cheng, Shiliang Pu, Yi Niu, Fei Wu, and Futai Zou. Spin: Structure-preserving inner offset network for scene text recognition. In *AAAI*, volume 35, pages 3305–3314, 2021. 7
- [52] Xinyun Zhang, Binwu Zhu, Xufeng Yao, Qi Sun, Ruiyu Li, and Bei Yu. Context-based contrastive learning for scene text recognition. In *AAAI*, pages 3353–3361, 2022. 6
- [53] Wenqi Zhao and Liangcai Gao. Comer: Modeling coverage for transformer-based handwritten mathematical expression recognition. In *ECCV*, volume 13688, pages 392–408, 2022. 1
- [54] Dajian Zhong, Shujing Lyu, Palaiahnakote Shivakumara, Bing Yin, Jijia Wu, Umapada Pal, and Yue Lu. Sgbnet: Semantic gan and balanced attention network for arbitrarily oriented scene text recognition. In *ECCV*, volume 13688, pages 464–480, 2022. 6