

Reading Chinese in Natural Scenes with a Bag-of-Radicals Prior

Yongbin Liu^{1,2}
liuyongbin@buaa.edu.cn

Qingjie Liu^{*1,2}
qingjie.liu@buaa.edu.cn

Jiaxin Chen¹
jiaxinchen@buaa.edu.cn

Yunhong Wang^{1,2}
yhwang@buaa.edu.cn

¹ State Key Laboratory of Virtual Reality
Technology and System
Beihang University
Beijing, China

² Hangzhou Innovation
Institute of Beihang University
Hangzhou, China

Abstract

Scene text recognition (STR) on Latin datasets has been extensively studied in recent years, and state-of-the-art (SOTA) models often reach high accuracy. However, the performance on non-Latin transcripts, such as Chinese, is not satisfactory. In this paper, we collect six open-source Chinese STR datasets and evaluate a series of classic methods performing well on Latin datasets, finding a significant performance drop. To improve the performance on Chinese datasets, we propose a novel radical-embedding (RE) representation to utilize the ideographic descriptions of Chinese characters. The ideographic descriptions of Chinese characters are firstly converted to bags of radicals and then fused with learnable character embeddings by a character-vector-fusion-module (CVFM). In addition, we utilize a bag of radicals as supervision signals for multi-task training to improve the ideographic structure perception of our model. Experiments show performance of the model with RE + CVFM + multi-task training is superior compared with the baseline on six Chinese STR datasets.

1 Introduction

The scene text recognition (STR) task is to recognize text in natural scenes, such as road signs, shop banners, house numbers, etc, making it attractive for many downstream applications, for instance, mobile language translators, autonomous driving and geo-location. Unlike traditional optical character recognition (OCR) that mainly aims at document recognition, the shapes of text in natural scene images are arbitrary, such as curved shapes; and more challenging are that they may be captured under low illumination, various distances, and fonts are diverse.

In recent years, extensive efforts have been devoted to scene text recognition [1, 2, 3, 4, 5]. These methods mainly focus on Latin, which are trained and tested on Latin datasets. Few works have considered recognizing east Asian texts such as Chinese, Japanese,

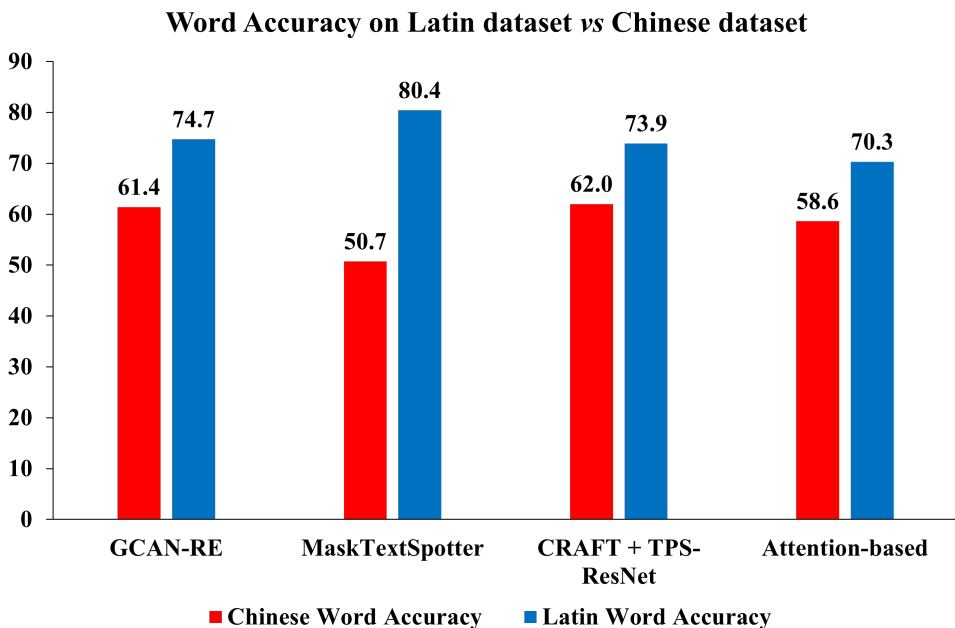


Figure 1: The performance of top-4 methods on Latin part and Chinese part of ICDAR-2019 Robust Reading Challenge on Arbitrary-Shaped Text leaderboard (scene text recognition task). There is a significant performance gap between Latin and Chinese recognition in natural scenes.

and Korean, which have distinct structures to Latin characters, and nearly 20% of world population are using them.

Fig. 1 shows the word accuracy on Latin transcripts and Chinese transcripts on ICDAR-2019 Robust Reading Challenge on Arbitrary-Shaped Text leaderboard [5], and it is obvious that there is a significant performance gap between Chinese part and Latin part. The word accuracy on Latin transcripts reach 70+%, while on Chinese transcripts is only approximately 60%, which is far lower.

To facilitate research in text recognition in natural scene images, ICDAR organized several competitions on multi-language scene text reading, including LSVT [17], ArT [6] and MLT [14]. These tracks release datasets that contain Chinese transcripts. However, Chinese samples are mixed together with other languages, making it difficult to assess performance on Chinese transcripts. Furthermore, the methods in competitions use different configurations of experiments and comparisons between them are unfair. In order to make a fair comparison between different approaches and benchmark Chinese character recognition in natural scenes, we collect several widely used Chinese STR datasets and evaluate several approaches on them with the same configuration. Since Latin transcripts are also common used in Chinese Street views, such as ‘SPA’ (balneotherapy), ‘KTV’ (Karaoke TV) and numbers on road signs, we evaluate the models on both Latin part and Chinese part, separately.

A model that performs well on Latin transcripts may have low performance on Chinese transcripts, indicating that we need to develop new models to read Chinese characters. Inspired by the learning process of pupils on Chinese characters, we utilize the radical infor-

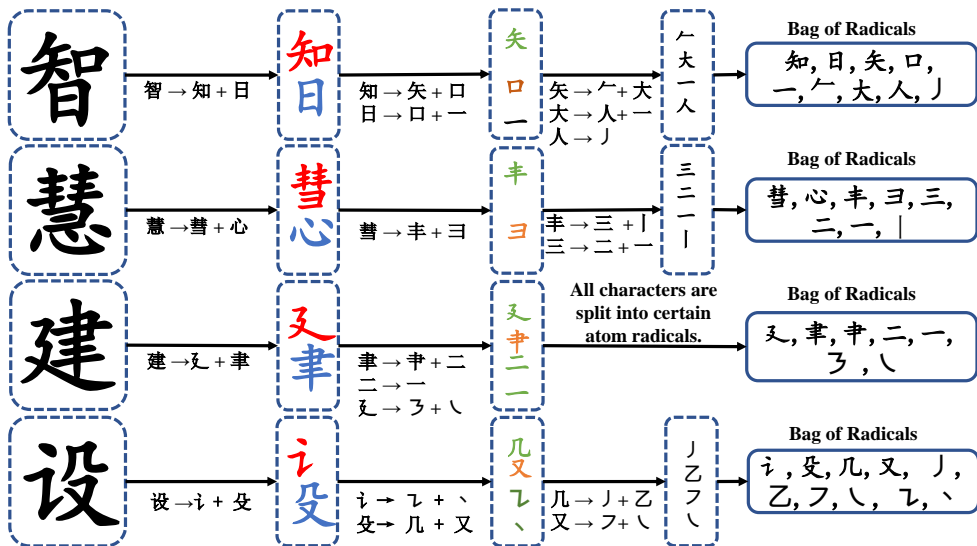


Figure 2: This figure demonstrates how we transform a Chinese character to a bag of radicals. Firstly, we split the character into several basic Chinese radicals, then we collect all radicals of this character to a bag of radicals.

mation of Chinese characters to improve the performance of Chinese STR.

The main contributions of this work are three-fold:

1. We evaluate 10 well-known STR models on six large-scale Chinese STR datasets. This result provides a fair comparison in Chinese STR field.
2. We find a performance gap between the Chinese recognition and the Latin recognition. And even worse, the models performing very well on Latin datasets may have a low performance on Chinese datasets. It indicates that simply deploying a Latin STR method in Chinese scenes may cause a performance drop.
3. A novel Chinese STR method that fuses the bags of radicals with CVFM is proposed, and we design a multi-task radical classification branch to improve the character structural perception. Experiments demonstrate that our method is superior to the baseline model.

2 Related Work

2.1 Scene Text Recognition

Early STR methods recognize words in a bottom-up manner [27]. With the development of deep learning techniques, researchers resort to convolutional neural networks (CNNs) as feature extractors and update them in a data driven manner. The optimized visual features are then reformulated by recurrent neural networks (RNN) as sequential features. The CNN and RNN are commonly defined as the encoder, transforming the original images into hidden representations. There are two families of decoders that convert the features into probabilistic

logits, that is CTC decoder [1, 2, 3] and 1D attention-based decoder [4, 5, 6]. Both CTC and 1D attention decoders frame the classification process as seq2seq learning like natural language processing (NLP) tasks.

Observing that the 1D attention decoder misses pixel-wise information and attention, [7] proposes a novel 2D attention decoder (SAR), which significantly improves the recognition performance on Latin datasets. Gaussian constrained attention networks (GCAN) [8] utilizes a learnable gaussian constrained module to improve the qualities of attention masks.

2.2 Chinese Character Recognition with Radicals as Priors

Several works have noticed the use of Chinese character radicals as a kind of knowledge prior. [9] and [10] propose DenseRAN, a single-character recognition network that firstly predicts the structure code of a Chinese character image and then transforms the code to a Chinese character. The decoding mechanism of [9] and [10] is designed for *single-character* recognition, and for a transcript that contains several Chinese characters, the radical code would be extremely long, resulting in impracticable word accuracy. [11] uses Chinese radicals as a kind of structural prototypes for few shot character recognition. Similar to DenseRAN, [11] also decomposes Chinese characters into basic radicals and encodes a character image into a radical sequence. [12] proposes to recognize Chinese characters by multi-label classification. [12] mainly focuses on printed single characters, and cannot be directly used to recognize scene text words. [13] proposes a stroke-based priors for Chinese character recognition. It first uses visual recognizer to encode the characters into stroke sequences, and then lookup from a lexicon for the most similar character as the prediction. Actually it is not a radical-based method since a radical consists several basic strokes.

Most above methods are designed for *single-character* recognition and only conduct experiments on *single-character* datasets. Their application in scene text recognition is intractable due to the long radical coding. In our method, we take bags-of-radicals for two usages, the first one is to use it as a embedding feature to fuse with visual features; the other is using them as supervision signals to build up a structure-awareness branch. The usage of radicals of our method is novel and different from [13], [9], [10], [11], [12], [12], [13], [14], [15], [16], [17] and [18].

3 The Proposed Approach

In this section, we describe details about the framework and the key components of the proposed method.

3.1 Framework Overview

Our method follows the same pipeline as GCAN [8], while the main difference lies in the decoder part. As illustrated in Fig. 4, given an input image, the 2D visual feature map V is extracted by a ResNet backbone. h_w is the contextual information encoded by an LSTM cell. The decoder predicts the characters in a sequential step-by-step manner, where h'_t indicates the contextual information. Specifically, at time step T , the precedent prediction y_{T-1} is firstly fed into the proposed Character Vector Fusion Module (CVFM) and transformed into a character representation F_{RE}^{T-1} with the Bag-of-Radicals prior, which is subsequently taken as the input of an LSTM cell to generate the contextual knowledge h'_T . h'_T is there-

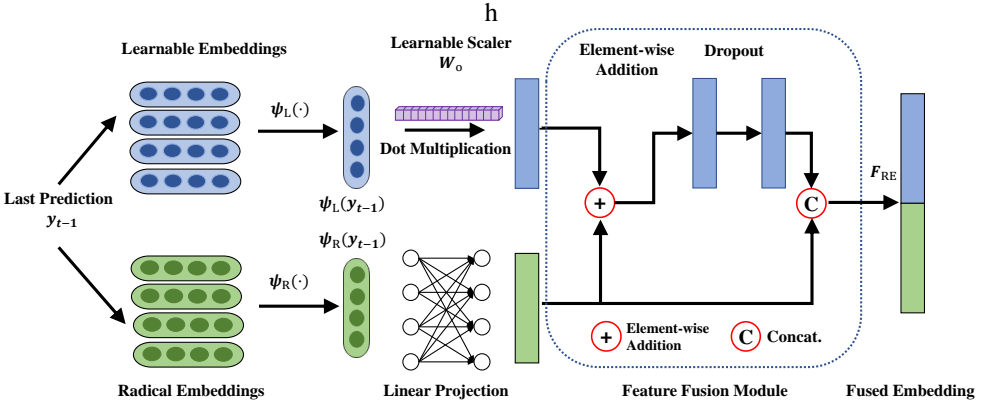


Figure 3: The pipeline of our character vector fusion module. Firstly the looked up learnable embedding of y_{t-1} is scaled by a learnable weight \mathbf{W}_o to $\mathbf{W}_o \cdot \mathbf{o}$, and the radical embedding is linearly transformed by \mathbf{W}_r to $\mathbf{W}_r \mathbf{r}$, which is projected into the same vector space as $\mathbf{W}_o \cdot \mathbf{o}$. These two features are summed and sent to a dropout layer, which aims to protect the features from overfitting. The result is concatenated with radical embedding \mathbf{r} , and that is the fused embedding feature \mathbf{F}_{RE} .

after integrated with the visual map \mathbf{V} by passing a 2D attention module \textcircled{A} with an output \mathbf{g}_T . \mathbf{g}_T is further fused with \mathbf{h}'_T by a Gaussian Constrained Refinement Module (GRCM), yielding a refined vector \mathbf{g}'_T , which is ensembled with \mathbf{g}_T by element-wise addition. The final representation \mathbf{F}_T is used to predict the character \mathbf{y}_T via a predictor $\phi(\cdot)$. To exploit the radical prior for learning a discriminative \mathbf{F}_T , we employ an auxiliary task, *i.e.* predicting which radicals exist in the current character, by adopting a radical classification head $\phi_R(\cdot)$. Accordingly, a radical-induced hybrid loss is developed to optimize the overall network.

3.2 Character Vector Fusion Module

In this section, we will describe the CVFM module. Before that, we firstly employ the Bag-of-Radicals Embedding for an arbitray Chinese character.

Bag-of-Radicals Embedding. Given a Chinese character, we utilize the open source Unicode dictionary *cjkvi* [4] to progressively decompose it into a bag of radicals. For instance, as shown in Fig. 2, the character ‘慧’ can be firstly divided into two principle parts, *i.e.*, ‘彗’ and ‘心’. The radical ‘彗’ can be further separated into ‘彗’ and ‘丰’. The decomposition process as above continues until the ‘丰’ is decomposed into *atom* radicals ‘三’, ‘二’, ‘一’ and ‘丨’. By this mean, an arbitrary Chinese character can be represented by a bag of radicals as well as its atom radicals. For instance, the bag-of-radicals representation of ‘慧’ is {彗, 心, 丰, 彗, 三, 二, 一, 丨}

By performing the bag-of-radicals (BoR) decomposition on all characters in the Chinese charset (totally 5,941 characters), we can obtain a bank of 1,637 distinct basic Chinese radicals $CRBank = \{CR_1, CR_2, CR_3, \dots, CR_D\}$, where $D=1637$. Based on $CRBank$, given an arbitrary Chinese character C , we encode it into a 1,637 dimensional BoR vector $\mathbf{r}_C = [v_1, v_2, \dots, v_{1637}] \in \mathbb{R}^{1637}$ according to its BoR decomposition $\{CR_{i_1}, CR_{i_2}, \dots, CR_{i_K}\}$, where $v_j = 1$ for $j = i_1, i_2, \dots, i_K$, and 0 otherwise. As for any Latin character C_L that does not contain Chinese radicals, we set the BoR of it a zero vector. By encoding all 5,941

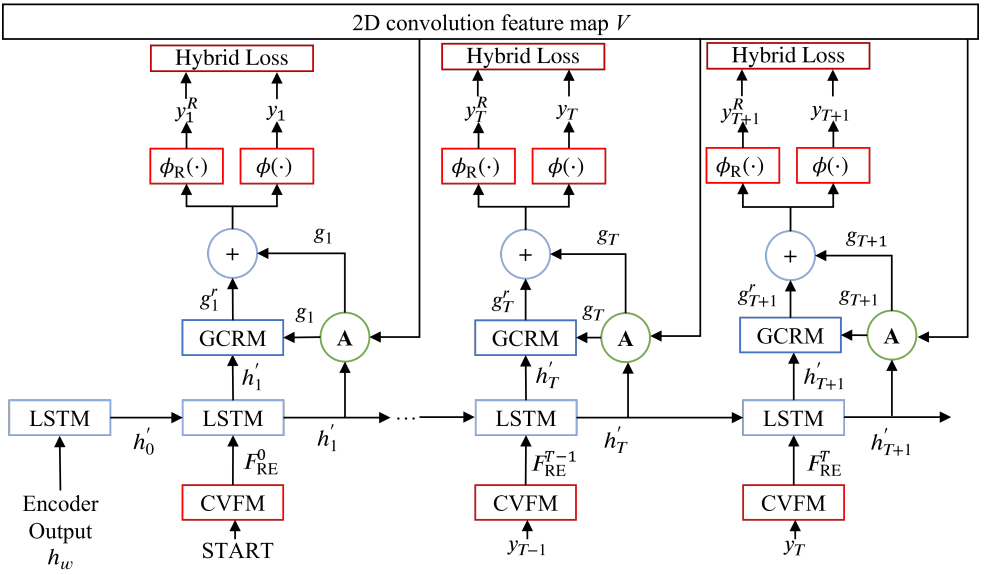


Figure 4: The structure of the decoder used in our model. The contextual feature h'_t , visual feature map extracted by backbone networks V and the embedding vector of last predicted character $\psi(y_{t-1})$ are fused by 2D attention modules. The GCRM module is to use Gaussian masks [15] to constrain the attention map. \textcircled{A} is 2D attention, and \oplus means element-wise addition.

characters into the BoR vectors, we can obtain a matrix $\mathbf{R} = [\mathbf{r}_1; \dots; \mathbf{r}_{5941}] \in \mathbb{R}^{5941 \times 1637}$.

CVFVM. We now introduce the character vector fusion module (CVFVM) which integrates the geographic information. It first converts the predicted character at the last time step y_{T-1} to radical embeddings (RE) and then fuses RE with a learnable embedding (LE) to obtain fusion \mathbf{F}_{RE} . The decoder uses the fused features to compute contextual information by an LSTM cell, as shown in Fig. 4, and further makes prediction. The fusion steps are described as follows:

$$\mathbf{o} = \psi_L(\mathbf{y}_{t-1}) \quad (1)$$

$$\mathbf{r} = \psi_R(\mathbf{y}_{t-1}) \quad (2)$$

$$\mathbf{F}_{RE} = \text{Concat}(\text{Dropout}(\mathbf{W}_r \mathbf{r} + \mathbf{W}_o \cdot \mathbf{o}); \mathbf{r}) \quad (3)$$

Firstly, the looked up learnable embedding of \mathbf{y}_{t-1} is scaled by weight \mathbf{W}_o to obtain $\mathbf{W}_o \cdot \mathbf{o}$, and the radical embedding is linearly transformed by \mathbf{W}_r to the same dimensionality as \mathbf{o} . A dropout layer is applied to avoid overfitting. The result is concatenated with radical embedding \mathbf{r} , finally obtaining the fused embedding feature \mathbf{F}_{RE} . This mechanism is shown in Fig. 3. We use the linear projection to transform the radical embedding into the same semantic space as learnable embedding, and the dot multiplication is only to adjust the scale of learnable embedding.

In the baseline (GCAN), the decoder only uses a learnable embedding to extract contextual knowledge, which does not contain any structural prior. Our modification is to utilize the radical prior to boost the learning of contextual features.

3.3 A Radical-Induced Hybrid Loss

In addition to the use of radical embeddings, we also build a multi-task branch to boost the character-structure awareness of our model. In the prediction head of the decoder, we use the bags-of-radicals as ground truth to supervise the model. The hybrid loss in our model is defined as follows:

$$\hat{\mathbf{y}}_T = \phi(\mathbf{g}_T) \quad (4)$$

$$\hat{\mathbf{y}}_T^R = \phi_R(\mathbf{g}_T) \quad (5)$$

$$\mathcal{L}_O = \text{CrossEntropy}(\hat{\mathbf{y}}_T, \mathbf{y}_T) \quad (6)$$

$$\mathcal{L}_R = \text{BCE}(\hat{\mathbf{y}}_T^R, \mathbf{y}_T^R) \quad (7)$$

$$\mathcal{L} = \mathcal{L}_O + \mathcal{L}_R \quad (8)$$

where \mathbf{g}_T represents the glimpse vector in Fig. 4. $\hat{\mathbf{y}}_T$ represents the probability logits of the prediction head of the baseline, and the corresponding ground truth \mathbf{y}_T is a one-hot vector. ϕ_R is the radical prediction classifier, and the output is $\hat{\mathbf{y}}_T^R$. The ground truth of the radical branch \mathbf{y}_T^R is the proposed radical embeddings. We use cross entropy loss (CE loss) and binary cross-entropy loss (BCE loss) to optimize the two branches respectively, and the losses are summed up to the hybrid loss. It is worth noting that the radical representations of Latin characters are null and not taken into BCE loss computation. Specifically, we use a mask to label zero for the Latin characters and one for Chinese characters, and the mask are multiplied by the loss matrix, and we demonstrate how this mask works in supplementary materials.

4 Experiments

In this section, we conduct experiments to show the superiority of using radical embeddings, CVFM and the proposed hybrid loss for reading Chinese in natural scenes. We compare our model with baselines and conduct performance analysis on the evaluation results. We also utilize ablation study to verify each component of our model. The *case study*, *performance analysis on Latin subset* and *compatibility with text detectors* are demonstrated in supplementary materials.

4.1 Dataset Details

Table 1: The number of train and test samples in each Chinese dataset.

Dataset	Year	Train		Test	
		Total	Non-Latin	Total	Non-Latin
ArT	2019	39,493	11,280	9,748	2,798
CASIA	2018	57,599	41,091	25,097	17,889
ctw	2018	176,472	176,448	12,562	12,560
LSVT	2019	213,985	176,814	24,143	20,097
RCTW	2017	34,727	26,901	9,017	6,990
ReCTS	2019	86,755	67,697	22,033	17,079
Total		609,031	500,231	102,600	77,413

ICDAR-2019 ArT(ArT) is proposed in [5] to evaluate the STR models on texts with complex shapes, such as curved shapes. There are 49,241 cropped images in the training dataset. Since the test dataset does not provide labels, we randomly split the training set into two parts. 20% (9,748 items) of them is used as testing, and 80% (39,493 items) serves as the training set. Among all characters, 25.8% are Chinese and the rest of 74.2% are Latin.

CASIA-10k (CASIA) is proposed by [6] for evaluating scene text detections. Bounding boxes are annotated with text content and cropped into samples. Finally, there are 57,599 samples in the training set and 25,097 in the test set.

Large Chinese Text Dataset in the Wild (ctw) is proposed by [7], mainly containing street view images captured by cameras. Based on bounding boxes, there are 176,472 cropped images for training, and 12,562 for test.

Large-scale Street View Text (LSVT) is a large-scale street view text dataset [8]. There are 30k images which contain bounding boxes and transcripts. 213,985 cropped images are used in the training part and 24,143 are used for the test part.

ICDAR2017 Competition on Reading Chinese Text in the Wild (RCTW) [9] is created for the ICDAR-2017 competition. Since the ground truth files of test datasets are not publicly available, we firstly crop the images in the training dataset and obtain 43,744 samples. 20% (9,017) are randomly selected to compose our test dataset and the remaining 34,727 samples for training.

Reading Chinese Text on Signboard (ReCTS) is created by [5] for ICDAR 2019 STR. Each image is annotated with bounding boxes of rectangles. Since the test dataset is not publicly available, we randomly divide the training dataset into two parts. 86,755 samples (80%) are used as training samples and the other 22,033 images are for evaluation.

We also demonstrate the number of non-Latin samples in each dataset in Tab. 1 (*non-latin* samples contain at least one Chinese character).

4.2 Experimental Setups

We build up the STR model pipeline using MMOCR [9] and deep STR benchmark [10]. To make fair comparisons, all models adopt the same experimental setting and are trained on the same training data, that is ArT + CASIA + ctw + LSVT + RCTW + ReCTS (609,031 samples in total, as shown in Tab. 1). The evaluation datasets are the test part of above six datasets (102,600 samples in total). During training, we adopt the Adam optimizer [8], with an initial learning rate 0.001 and no weight decay. We use pretrained parameters on synthetic Latin datasets to initialize the model. The batch size during training is 1024, distributed on eight GTX 2080Ti GPUs. The height of each input image is rescaled to 128 pixels while keeping the height/weight ratio invariant; if the new width is greater than 256 pixels, we resize the image to 128×256 , otherwise pad it to 128×256 with zeros. Data augmentation is applied during training, including color jitter transformation (brightness 0.5, contrast 0.5, saturation 0.5, hue 0.5), the random affine transformation (degrees 0.5, translate (0.1, 0.1), scale (0.95, 1)), the random rotation transformation (degrees 5), and the identity transformation. For each image sample, we randomly select one of the above augmentation types in each mini-batch.

4.3 Evaluation Protocols

The common metric to evaluate the similarity of two transcripts is word accuracy (two strings are considered the same if and only if every two aligned characters in the string pair are the same). In order to eliminate the influence of symbols and spaces, previous STR works ignore all symbols (!@#%\$, etc), and we follow the practice in this work.

Another measure in ICDAR competitions [6] is to use 1 - Normalized Edit Distance(1-N.E.D) to evaluate the similarity of two strings. The computation is $1 - \frac{\text{EditDis}(s_1, s_2)}{\max(\text{len}(s_1), \text{len}(s_2))}$. In our experimental results, we report model performance using both measures.

4.4 Experimental Results

Table 2: We evaluate models including CRNN [12], RARE [18], R2AM [10], STAR-NET [12], GRCNN [22], Rosetta [2], TRBA [4], SAR [10], RobScan [29], GCAN [16] on all six datasets using word accuracy and 1-N.E.D metrics.

Model	Word accuracy on all transcripts							1-Normalized edit distance on all transcripts						
	ArT	CASIA	ctw	LSVT	RCTW	ReCTS	All	ArT	CASIA	ctw	LSVT	RCTW	ReCTS	All
GRCNN	42.1	40.5	31.1	37.5	41.5	53.9	41.7	72.5	69.5	61.3	64.8	67.5	72.6	68.2
R2AM	60.1	53.6	53.2	58.7	55.5	67.6	58.5	78.9	74.9	69.8	75.5	73.9	80.6	75.9
CRNN	44.4	44.9	38.5	45.7	45.8	60.9	47.8	75.0	72.9	65.3	70.6	70.9	78.3	72.6
Rosetta	64.6	59.0	54.8	62.2	61.3	74.6	63.3	84.5	81.6	76.8	81.3	81.1	86.9	82.3
RARE	72.1	64.6	62.8	68.0	66.6	76.9	68.7	86.5	82.6	78.2	82.7	81.9	86.7	83.3
STAR-NET	59.5	52.9	49.5	54.6	56.1	71.0	57.7	80.6	76.4	69.0	74.1	75.9	84.1	77.0
TRBA	75.6	66.0	63.8	69.3	68.1	78.7	70.3	88.9	83.3	78.7	83.4	82.9	88.0	84.3
RobScan	78.1	68.8	66.8	72.1	70.3	80.6	72.9	89.9	85.8	82.4	86.0	85.0	89.4	86.5
SAR	74.8	69.0	65.2	71.3	69.9	80.8	72.2	88.2	85.5	81.6	85.3	84.5	89.5	86.0
GCAN	75.4	70.0	65.2	72.1	70.8	81.4	72.8	88.6	86.2	81.5	85.8	84.9	89.8	86.4
Ours	78.2	71.3	67.5	73.9	72.5	82.3	74.6	89.8	86.5	82.5	86.5	85.6	90.2	87.0

We select 10 well known STR models to conduct experiments, including CRNN [12], RARE [18], R2AM [10], STAR-NET [12], GRCNN [22], Rosetta [2], TRBA [4], SAR [10], RobScan [29] and GCAN [16]. These models can be divided into three categories according to decoders they use: CTC-based (CRNN, GRCNN, Rosetta, STAR-NET), 1D attention-based (R2AM, RARE, TRBA), and 2D attention-based (SAR, RobScan, GCAN). The experimental results are shown in Tab. 2. Based on the experimental results, we reach the following conclusions:

- The models that adopt 2D attention decoders outperform models with other types of decoders. The lowest average word accuracy of 2D attention models on non-latin transcripts is 71.3, achieved by SAR. It is higher than any other methods using 1D attention decoders or CTC decoders such as TRBA (68.5) and Rosetta (63.7).
- Our method outperforms other models significantly. As evaluated on all transcripts, our method outperforms the baseline GCAN by 1.8% word accuracy and the largest margins are ArT (2.8%) and ctw (2.3%), indicating that our method works on both curved text and street view text images.

4.5 Ablation Study

In order to verify the function of our modules, we conduct an ablation study. We first remove the radical binary cross-entropy loss (BCE) and then the CVFM module. As shown in Tab. 3, using radical embeddings and fusion modules can improve the performance by 1.5% on non-Latin transcripts. The additional BCE loss can improve the performance by another 0.3%. The experimental results validate the effectiveness of CVFM and BCE loss.

Table 3: Ablation study w.r.t CVFM and BCE loss of our method. The first line shows the performance of our baseline (GCAN). The radical embeddings (RE) and CVFM are then added and the performance is shown in the 2nd line. The 3rd line demonstrates the effectiveness of our BCE Loss.

Ablation		Word accuracy on all transcripts							Word accuracy on non-Latin transcripts								
CVFM	BCE	ArT	CASIA	ctw	LSVT	RCTW	ReCTS	All	Gain	ArT	CASIA	ctw	LSVT	RCTW	ReCTS	All	Gain
✓		75.4	70.0	65.2	72.1	70.8	81.4	72.8	/	74.7	68.1	65.2	71.6	70.7	80.5	71.8	/
✓		77.4	71.0	67.6	73.6	72.2	82.1	74.3	+1.5	78.1	69.0	67.6	73.3	72.0	81.2	73.2	+1.4
✓	✓	78.2	71.3	67.5	73.9	72.5	82.3	74.6	+1.8	79.2	69.3	67.5	73.6	72.5	81.5	73.5	+1.7

We also conduct ablation study w.r.t each component of the CVFM module. We remove the scaler, dropout operation and set radical embedding to zero respectively. The experimental results are shown in Tab. 4. According to the results, removing any component or set the radical embedding causes significant performance drop, hence we validate the necessity of each component in CVFM.

Table 4: Ablation w.r.t components of the CVFM.

Ablation for CVFM	Word Accuracy on All Transcripts							Word Accuracy on non-Latin Transcripts						
	ArT	CASIA	ctw	LSVT	RCTW	ReCTS	All	ArT	CASIA	ctw	LSVT	RCTW	ReCTS	All
full CVFM	77.4	71.0	67.6	73.6	72.2	82.1	74.3	78.1	69.0	67.6	73.3	72.0	81.2	73.2
w/o scaler	70.4	66.1	65.2	69.6	68.1	79.7	70.3	71.3	64.7	65.2	69.3	68.3	79.0	69.7
w/o dropout	76.6	69.4	66.4	72.3	70.9	81.1	73.0	77.2	67.0	66.4	71.9	70.7	80.1	71.8
RE set zero	72.9	69.3	67.6	72.4	71.0	82.1	73.0	72.8	67.7	67.6	72.1	71.1	81.5	72.4

5 Conclusion

In this work we collect six Chinese scene text datasets to evaluate the previous models performing well on Latin datasets fairly. From the experimental results we can conclude that the 2D-attention models reach high word accuracy on both Latin datasets and Chinese datasets. In order to further improve the performance of Chinese STR, we propose RE, CVFM and hybrid radical losses for multi-task training. The experiments on the Chinese STR benchmark show that our method is superior compared with our baseline (GCAN) and other methods on six datasets.

References

- [1] Jeonghun Baek, Geewook Kim, et al. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *ICCV*, pages 4715–4723, 2019.
- [2] Fedor Borisyuk, Albert Gordo, et al. Rosetta: Large scale system for text detection and recognition in images. In *KDD*, pages 71–79, 2018.
- [3] Junbum Cha, Sanghyuk Chun, et al. Few-shot compositional font generation with dual memory. In *ECCV*, pages 735–751. Springer, 2020.
- [4] Jingye Chen, Bin Li, and Xiangyang Xue. Zero-shot chinese character recognition with stroke-level decomposition. 2021.

- [5] Chee Kheng Chng, Yuliang Liu, et al. Icdar2019 robust reading challenge on arbitrary-shaped text-rrc-art. In *ICDAR*, pages 1571–1576. IEEE, 2019.
- [6] Wenhao He, Xu-Yao Zhang, et al. Multi-oriented and multi-lingual scene text detection with direct regression. *TIP*, 27(11):5406–5419, 2018.
- [7] kawabata. Ids data. <https://github.com/cjkvi/cjkvi-ids>, 2018.
- [8] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [9] Zhanghui Kuang, Hongbin Sun, et al. Mmocr: A comprehensive toolbox for text detection, recognition and understanding. *arXiv preprint arXiv:2108.06543*, 2021.
- [10] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *CVPR*, pages 2231–2239, 2016.
- [11] Hui Li, Peng Wang, et al. Show, attend and read: A simple and strong baseline for irregular text recognition. In *AAAI*, volume 33, pages 8610–8617, 2019.
- [12] Wei Liu, Chaofeng Chen, et al. Star-net: a spatial attention residue network for scene text recognition. In *BMVC*, volume 2, page 7, 2016.
- [13] Long-Long Ma and Cheng-Lin Liu. A new radical-based approach to online handwritten chinese character recognition. In *2008 19th International Conference on Pattern Recognition*, pages 1–4. IEEE, 2008.
- [14] Nibal Nayef, Yash Patel, et al. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition—rrc-mlt-2019. In *ICDAR*, pages 1582–1587. IEEE, 2019.
- [15] Song Park, Sanghyuk Chun, Junbum Cha, Bado Lee, and Hyunjung Shim. Few-shot font generation with localized style representations and factorization. In *AAAI*, 2021.
- [16] Zhi Qiao, Xugong Qin, et al. Gaussian constrained attention network for scene text recognition. In *ICPR*, pages 3328–3335. IEEE, 2021.
- [17] Baoguang Shi, Xiang Bai, et al. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *TPAMI*, 39(11):2298–2304, 2016.
- [18] Baoguang Shi, Xinggang Wang, et al. Robust scene text recognition with automatic rectification. In *CVPR*, pages 4168–4176, 2016.
- [19] Baoguang Shi, Cong Yao, et al. Icdar2017 competition on reading chinese text in the wild (rctw-17). In *ICDAR*, volume 1, pages 1429–1434. IEEE, 2017.
- [20] Yipeng Sun, Zihan Ni, et al. Icdar 2019 competition on large-scale street view text with partial labeling-rrc-lsvt. In *ICDAR*, pages 1557–1562. IEEE, 2019.
- [21] An-Bang Wang and Kuo-Chin Fan. Optical recognition of handwritten chinese characters by hierarchical radical matching method. *Pattern Recognition*, 34(1):15–35, 2001.

- [22] Jianfeng Wang and Xiaolin Hu. Gated recurrent convolution neural network for ocr. In *NeurIPS*, pages 334–343, 2017.
- [23] Tianwei Wang, Zecheng Xie, Zhe Li, Lianwen Jin, and Xiang Chen. Radical aggregation network for few-shot offline handwritten chinese character recognition. *Pattern Recognition Letters*, 125:821–827, 2019.
- [24] Tie-Qiang Wang, Fei Yin, and Cheng-Lin Liu. Radical-based chinese character recognition via multi-labeled learning of deep residual networks. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 579–584. IEEE, 2017.
- [25] Wenchao Wang, Jianshu Zhang, Jun Du, Zi-Rui Wang, and Yixing Zhu. Denscran for offline handwritten chinese character recognition. In *2018 16th International Conference on Frontiers in Handwriting Recognition (ICFHR)*, pages 104–109. IEEE, 2018.
- [26] Changjie Wu, Zi-Rui Wang, Jun Du, Jianshu Zhang, and Jiaming Wang. Joint spatial and radical analysis network for distorted chinese character recognition. In *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW)*, volume 5, pages 122–127. IEEE, 2019.
- [27] Cong Yao, Xiang Bai, et al. Strokelets: A learned multi-scale representation for scene text recognition. In *CVPR*, pages 4042–4049, 2014.
- [28] Tai-Ling Yuan, Zhe Zhu, et al. A large chinese text dataset in the wild. *Journal of Computer Science and Technology*, 34(3):509–521, 2019.
- [29] Xiaoyu Yue, Zhanghui Kuang, et al. Robustscanner: Dynamically enhancing positional clues for robust text recognition. In *ECCV*, pages 135–151. Springer, 2020.
- [30] Jianshu Zhang, Yixing Zhu, et al. Radical analysis network for zero-shot learning in printed chinese character recognition. In *ICME*, pages 1–6. IEEE, 2018.
- [31] Rui Zhang, Yongsheng Zhou, et al. Icdar 2019 robust reading challenge on reading chinese text on signboard. In *ICDAR*, pages 1577–1581. IEEE, 2019.