

# RZCR: Zero-shot Character Recognition via Radical-based Reasoning

Xiaolei Diao<sup>1,2</sup>, Daqian Shi<sup>2</sup>, Hao Tang<sup>3</sup>, Qiang Shen<sup>1</sup>, Yanzeng Li<sup>4</sup>, Lei Wu<sup>4</sup>, Hao Xu<sup>1\*</sup>

<sup>1</sup>College of Computer Science and Technology, Jilin University

<sup>2</sup>DISI, University of Trento

<sup>3</sup>CVL, ETH Zurich

<sup>4</sup>College of Software Engineering, Jilin University

{xiaolei.diao, daqian.shi}@unitn.it, hao.tang@vision.ee.ethz.ch,

{shenqiang19, yzli20, wulei20}@mails.jlu.edu.cn, xuhao@jlu.edu.cn

## Abstract

The long-tail effect is a common issue that limits the performance of deep learning models on real-world datasets. Character image datasets are also affected by such unbalanced data distribution due to differences in character usage frequency. Thus, current character recognition methods are limited when applied in the real world, especially for the categories in the tail that lack training samples, e.g., uncommon characters. In this paper, we propose a zero-shot character recognition framework via radical-based reasoning, called RZCR, to improve the recognition performance of few-sample character categories in the tail. Specifically, we exploit radicals, the graphical units of characters, by decomposing and reconstructing characters according to orthography. RZCR consists of a visual semantic fusion-based radical information extractor (RIE) and a knowledge graph character reasoner (KGR). RIE aims to recognize candidate radicals and their possible structural relations from character images in parallel. The results are then fed into KGR to recognize the target character by reasoning with a knowledge graph. We validate our method on multiple datasets, and RZCR shows promising experimental results, especially on few-sample character datasets.

## 1 Introduction

Developments in optical character recognition (OCR) technology offer new solutions for learning, managing, and utilizing character resources. Current OCR methods are mainly based on deep learning models, which place high demands on the quantity and quality of data [Zhang *et al.*, 2017]. Due to differences in character usage frequency, the long-tail effect exists as a common issue in character datasets. Such datasets often contain categories with few samples, especially for unreproducible or inaccessible cases, e.g., historical character [Huang *et al.*, 2019] and calligraphic character [Lyu *et al.*, 2017] datasets. Fig. 1(a) shows the distribution of character samples for each category in an oracle bone dataset [Wu,

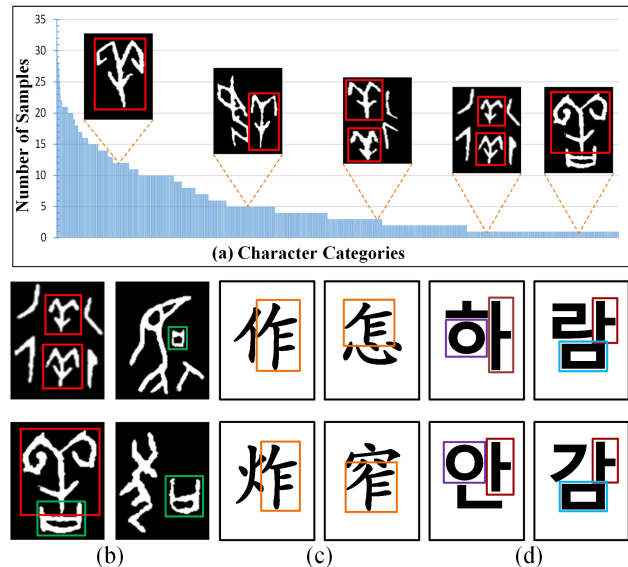


Figure 1: Analysis of characters, where radicals are distinguished by colored boxes. (a) Statistics of a character dataset; (b)-(d) Examples of oracle bone, Chinese, and Korean characters.

2012]. The distribution demonstrates a typical long-tail effect, where over half of the character categories have five or fewer samples, and there are even categories with a single sample. As a result, such datasets challenge the validity of current OCR methods.

Inspired by studies on general image classification, current research on character recognition mainly focused on methods based on deep convolutional neural networks (CNN). Such OCR methods achieve limited performance on unbalanced datasets, especially when recognizing character categories with few samples [Anderson, 2006; Cao *et al.*, 2020]. A common solution to alleviate the limitations caused by the few-sample problem is employing data augmentation methods to balance data categories, including re-sampling and re-weighting. However, these methods are rarely applied to data categories with few samples or one sample. Although the compromise solution tends to discard categories with few samples [Géron, 2019], we consider such categories valuable and irreplaceable for real-world tasks. For example, many

\*Corresponding Author

categories in Fig. 1(a) contain only one sample in the unearthed oracle bone dataset, which is unique evidence for understanding this character by archaeologists. Thus, we aim to achieve character recognition on few-sample datasets while keeping data integrity.

The orthography-based strategy of learning East Asian characters in terms of radicals and character structure is the most frequent and effective way for human beings [Shen, 2005]. We observe the orthographic peculiarities [Myers, 2016] of East Asian characters: (1) Characters are composed of a set of radicals organized in a specific structure, where radicals are semantic units of the character. (2) Radicals are commonly shared by different characters, thus, the number of radicals is less than that of characters significantly. Fig. 1(b), (c) and (d) show our observations on oracle bone, Chinese, and Korean characters, where we highlight radicals by colored boxes. Statistically, 6,763 Chinese characters can be represented by 485 radicals and 10 character structures [Liu *et al.*, 2011]. In Fig. 1(a), we demonstrate some characters that contain the same radical “yang” (in red boxes), where such characters tend to distribute randomly in the statistics. As a result, decomposing and reconstructing characters by exploiting radicals help to intentionally analyze characters, which offers the possibility to learn from enough radical samples and recognize character categories with few samples.

According to the above-mentioned discussions, we propose a novel method for zero-shot character recognition via radical-based reasoning, namely RZCR. Specifically, the proposed RZCR consists of a radical information extractor (RIE) and a knowledge graph-based character reasoner (KGR). The former aims to extract visual semantic information in characters to obtain candidate radicals and their structural relations in parallel through the radical attention blocks (RABs) and structural relation blocks (SRBs), respectively. A fused attention layer is proposed to fuse visual semantic information extracted by RABs and SRBs, and a dual spatial attention layer in RABs is introduced to handle the radical overlapping issues. The results of RIE are then fed into KGR, where a weight-fusion reasoning algorithm is proposed to reconstruct characters by modeling the radical information as soft labels. We achieve zero-shot recognition by reasoning with the pre-designed character knowledge graph (CKG) that stores information about characters, radicals, and their relations.

The contributions of our paper are summarized as follows:

- We propose a novel zero-shot method (i.e., RZCR) for character recognition, which can effectively handle categories with insufficient training samples by character decomposing and reconstructing.
- RZCR introduces a new strategy to organize candidates for reasoning-based character recognition, where candidate radicals and structural relations are extracted in parallel by RIE, and then characters are recognized by the proposed weight-fusion reasoning algorithm in KGR.
- Our method is validated on multiple datasets, including a newly constructed dataset, namely OracleRC. Compared to state-of-the-art OCR methods, our RZCR achieves promising results, especially on few-sample cases.

## 2 Related Work

**Zero-shot Learning.** Zero-shot learning aims to classify unseen categories by learning from existing categories and knowledge supplemented by auxiliary resources [Wang *et al.*, 2019b; Sun *et al.*, 2021]. In recent years, zero-shot learning has been considered from a variety of perspectives. For instance, embedding-based methods [Akata *et al.*, 2015; Xie *et al.*, 2019] propose to perform image embedding and label semantic embedding in the same space and extend to unseen categories by a compatibility function. Furthermore, attribute-based methods [Lampert *et al.*, 2009; Lampert *et al.*, 2013] manually design attributes in different ways and represent categories with multidimensional attribute vectors, which are introduced to train the visual classifier by vector mapping. Moreover, reasoning-based methods exploit knowledge graphs [Rohrbach *et al.*, 2011; Wang *et al.*, 2018b] to guide zero-shot learning since knowledge graphs connect trained categories and unseen categories by pre-defined relations. The success of the above methods brings us inspiration about the usage of auxiliary information from knowledge graphs and the organization of attributes in zero-shot character recognition.

**Character Recognition.** Early studies on character recognition mainly contain feature extraction-based [Su and Wang, 2003] and statistical-based methods [Shanthi and Duraiswamy, 2010], which perform poorly on datasets with a large number of categories. Then, researchers exploit CNN, e.g., [Yuan *et al.*, 2012] applies an improved LeNet-5 model to recognize English characters, and [Cireşan and Meier, 2015] first introduces CNN into Chinese character recognition. To recognize curved or distorted characters in the real world, [Zhan and Lu, 2019; Yang *et al.*, 2019] introduce trainable models based on symmetry-constrained rectification and line-fitting transformation, respectively. Moreover, deep character recognition models with dedicated improvements achieve promising results in different languages [Balaha *et al.*, 2021; Mushtaq *et al.*, 2021]. The success of the above deep learning-based methods relies on a large number of training samples for each character category. Thus, data augmentation strategies are introduced to solve insufficient data issues. For instance, [Qu *et al.*, 2018] combines global transformation and local distortion to effectively enlarge the training dataset, and [Luo *et al.*, 2020] designs a set of custom fiducial points to assist in flexibly enhancing character images, both of which alleviate the insufficient training sample issue. Some researchers also introduce domain knowledge related to character attributes to deal with few-sample problems by decomposing characters, i.e., presenting characters as preset sequences. [Cao *et al.*, 2020; Zhang *et al.*, 2020] introduce CNN-based encoder-decoder frameworks to generate radical sequences for character recognition. Inspired by these methods, [Chen *et al.*, 2021] decomposes characters into strokes, the smallest units of characters, then perform recognition by looking up generated stroke sequences in a dictionary. However, the above methods require the generated sequences to match the dictionary exactly, both for each element and the order, i.e., hard-matching strategy. As a result, such sequence-based methods perform poorly on the

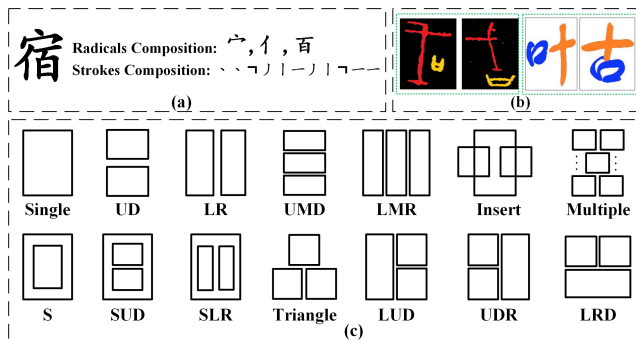


Figure 2: (a) Character composition in radical and stroke level, respectively; (b) Examples of characters composed by same radicals (highlighted in same colors) but different structural relations; (c) 14 predefined structural relations.

benchmarks, which limits their applications in practice.

### 3 Intuitive Discussion

In this section, we provide some intuitive observations to explore the character recognition task and discuss potential performance improvement strategies, driving the motivation of this paper. Then, we model the zero-shot character recognition to clarify the task in this paper.

**Decomposition of Characters.** According to orthography, characters can be decomposed into a set of radicals organized in specific structures, where radicals are independent characters or evolve from simple-semantic characters [Yeung *et al.*, 2016]. Complex-semantic characters usually contain more than one radical since radicals are the smallest units that represent complete semantic information [Ho *et al.*, 2003]. A Chinese character is shown in Fig. 2(a) as an example, where the three radicals in the first row represent “house”, “person”, and “mat (evolved)” and constitute the character that means “get an accommodation”. However, the decomposition in the second row shows an opposite example, i.e., the character is decomposed excessively into strokes, where the major semantic information is lost. Meanwhile, we define the structure of character as the relative localization relation between the constituent radicals, namely structural relations, which also contain semantic information. Structural relations are also crucial for characters since the same radical composition points to different candidate characters when the structure changes, as shown in Fig. 2(b). Moreover, as discussed earlier, the same radical can be shared by various characters, e.g., 2,374 Korean characters contain 68 radicals [Zatsepin *et al.*, 2019], and over 6,000 common Chinese characters can be composed of less than 500 radicals [Liu *et al.*, 2011], which means a large number of characters can be presented by a small number of radicals.

**Motivation and Challenges.** Based on the above observations, two ideas are derived for achieving zero-shot character recognition. Firstly, a character can be properly decomposed into a set of radicals and their structural relations, i.e., radical information, which can be extracted from the character image. Since the semantics of characters can be represented by radical information, we aim to extract such semantic infor-

Method	ASL	Accuracy Calculation	Example ( $r_i, s_i = 0.9$ )
Radical level	7.76	$p = \prod_{i=0}^m r_i$	$0.9^7 = 0.4783$
Stroke level	12.88	$p = \prod_{i=0}^n s_i$	$0.9^{12} = 0.2824$

Table 1: Accuracy estimation for sequence-based methods. ASL refers to the average sequence length.

mation through deep neural networks. Thus, we propose RIE to extract all the radical information in parallel rather than extracting separately since radicals and the structure support each other. Meanwhile, we apply the attention mechanism to deal with the radical overlapping problem for better recognition performance.

Secondly, current sequence-based methods organize character elements in the form of long sequences that are matched by hard-matching strategies. In Table 1, we show accuracy estimation for sequence-based methods at both radical and stroke levels, where  $p$  represents the probability of correct character recognition;  $r_i, s_i$  are correct recognition probabilities of elements in radical and stroke sequences, respectively;  $m, n$  are the lengths of sequences. We find  $p$  only achieves 0.48/0.28 maximally under the hard-matching strategy, even if we ideally assume a high recognition probability  $r_i=s_i=0.9$ . Therefore, we propose KGR to achieve reasoning-based zero-shot recognition by exploiting knowledge graphs to organize characters, radicals, and structural relations in a flexible way rather than in long sequences. Meanwhile, we introduce soft labels into the matching strategy, i.e., the proposed weight-fusion algorithm, to achieve stable recognition performance.

**Task Formulation.** The goal of the character recognition is to classify character image  $I_C \in \mathbb{R}^{H \times W \times C}$  into character category  $Z$ , where  $H \times W$  represents the spatial resolution, and  $C$  is the number of channels. We model character recognition as a zero-shot classification task, defined as follows. The training data is defined as  $I_{seen} = \{(x, y) | x \in X, y \in Y\}$ , where  $x$  is the input image from the training image set  $X$ ,  $y$  refers to the label(s) of  $x$  annotated based on radical information, and  $Y$  is the label set of the radical categories. Note that there are one or more radicals in an image. Similarly, the test set can be defined as  $I_{unseen} = \{(x, z) | x \in \hat{X}, z \in Z\}$ , where  $x$  belongs to images from unseen categories  $\hat{X}$ , i.e., character categories,  $Z$  represents the set of labels for unseen categories. In this task, the training is based on radical categories and finally outputs the target character category as  $f: x \rightarrow Z$ .

## 4 The Proposed RZCR

In this section, we detail the proposed RZCR and its key components RIE and KGR, as shown in Fig. 3.

### 4.1 Radical Information Extractor

The radicals and the structural relations in characters jointly convey character semantics, which shows both of them are indispensable for character recognition. RIE introduces two sets of blocks, Radical Attention Blocks (RABs) and Structural Relation Blocks (SRBs), aiming to obtain candidate rad-

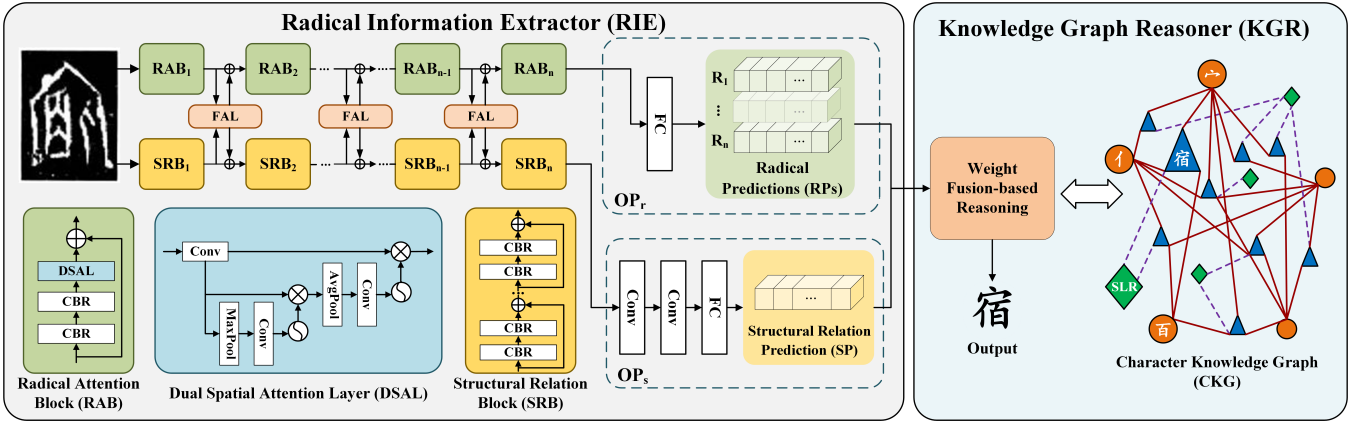


Figure 3: The overall architecture of RZCR. (left) RIE extracts candidate radicals and structural relations from input character images; (right) KGR recognizes the target character by reasoning with the CKG, where blue, orange, and green nodes represent characters, radicals, and structural relations, respectively, and lines in CKG represent different relations between nodes.

icals and structural relations from an input character image  $I_C$ , respectively. We observe that character images usually contain overlapping and unclear boundaries between radicals [Shi *et al.*, 2022b], especially for handwritten characters, which challenges the performance of radical recognition. Thus, the RABs apply the attention mechanism to concentrate on the target radicals. RAB consists of a dual spatial attention layer (DSAL) and two consecutive CBR blocks (including a  $3 \times 3$  convolutional layer, a batch normalization operation, and a ReLU operation), and SRB consists of a set of residual CBR blocks, as shown in the Fig. 3. For RAB, each DSAL obtains the attention weights by two calculations, in which the foreground information is retained for radical feature selection via MaxPool, and the background information is processed via AvgPool, to maintain integrity. The process in DSAL can be defined as:

$$\begin{aligned} A'_i &= \sigma(\text{Conv}(\text{MaxPool}(F_i))) \otimes F_i, \\ A_i &= \sigma(\text{Conv}(\text{AvgPool}(A'_i))) \otimes F_i, \end{aligned} \quad (1)$$

where  $A'_i$  and  $A_i$  denote the intermediate and final output features of the DSAL in the  $i$ -th RAB;  $\sigma$  refers to the sigmoid function; Conv refers to the convolutional layer;  $F_i$  is the feature map fed into the corresponding DSAL.

A fused attention layer (FAL) is proposed to fuse visual semantic information extracted by RABs and SRBs, aiming to improve the performance of radical prediction by injecting structural information and position features into RABs. In FAL, we introduce channel attention and spatial attention to provide efficient tokens from selective feature maps and to enhance dependency extraction on spatial axes, respectively, which is inspired by [Shi *et al.*, 2022a]. Given an input feature  $F_k$ , the output of fused attention layer  $FAL(F_k)$  is:

$$FAL(F_k) = M_s(M_c(F_k) + F_k) + M_c(F_k) + F_k, \quad (2)$$

where  $M_c$ ,  $M_s$  are channel and spatial attention, respectively. We denote the  $i$ -th feature  $F_{RA}$  as:

$$F_{RA_i} = FAL(F_{R_i} + F_{S_i}) + F_{R_i}, \quad (3)$$

The  $i$ -th structural relation feature  $F_{SR}$  is expressed as:

$$F_{SR_i} = FAL(F_{R_i} + F_{S_i}) + F_{S_i}, \quad (4)$$

Note that  $F_{RA_i}$  and  $F_{SR_i}$  will be sent to the next RAB and SRB, respectively. We have two output projectors  $OP_r$  and  $OP_s$  in RIE, which correspond to RAB and SRB, respectively. The  $OP_r$  consists of an FC layer whose size is  $K \times K \times M \times (n_r + n_c)$ , where  $K \times K$  refers to the number of divided grids of an input character image,  $M$  represents the number of anchor boxes in each grid,  $n_r$  is the number of radical categories in the datasets, and  $n_c$  records the coordinates of the radical location  $(x, y, w, h)$  and the confidence of radical detection, thus  $n_c = 5$ .  $OP_s$  consists of two convolutional layers and an FC layer to further handle features with mixed character semantics.

**Loss Functions.** RIE contains two output projectors  $OP_r$  and  $OP_s$ . Thus, we have corresponding loss functions  $\mathcal{L}_R$  and  $\mathcal{L}_S$ , respectively.  $\mathcal{L}_R$  consists of three components,  $\mathcal{L}_r$ ,  $\mathcal{L}_{coo}$  and  $\mathcal{L}_{isR}$ . We define  $\mathcal{L}_r$  as the cross-entropy loss of radical classification:

$$\mathcal{L}_r = \sum_{i=0}^{K^2} \sum_{j=0}^M \mathbb{I}_{i,j}^{isR} \sum_{r \in R} [p_i(r) \log(\hat{p}_i(r)) + (1 - p_i(r)) \log(1 - \hat{p}_i(r))], \quad (5)$$

where  $r$  is a category from radical set  $R$ ;  $\mathbb{I}_{i,j}^{isR}$  is assigned 1 or 0 to indicate whether the target radical exists in the proposal coordinates;  $p_i(r)$  and  $\hat{p}_i(r)$  refer to the probability of target and predicted results, respectively. We also define  $\mathcal{L}_{coo}$  as the loss function of the radical coordinates:

$$\begin{aligned} \mathcal{L}_{coo} &= \sum_{i=0}^{K^2} \sum_{j=0}^M \mathbb{I}_{i,j}^{isR} (2 - w_i h_i) [(x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2] \\ &+ \sum_{i=0}^{K^2} \sum_{j=0}^M \mathbb{I}_{i,j}^{isR} (2 - w_i h_i) [(w_i - \hat{w}_i)^2 + (h_i - \hat{h}_i)^2], \end{aligned} \quad (6)$$

$\mathcal{L}_{isR}$  is the loss function of the radical detection confidence:

$$\mathcal{L}_{isR} = \sum_{i=0}^{K^2} \sum_{j=0}^M \mathbb{I}_{i,j}^{isR} (c_i - \hat{c}_i)^2 + \lambda \sum_{i=0}^{K^2} \sum_{j=0}^M \mathbb{I}_{i,j}^{noR} (c_i - \hat{c}_i)^2, \quad (7)$$

where we consider both cases of radical existence and absence. Thus, we also set  $\mathbb{I}_{i,j}^{noR}$  to present if there is no radical

covered by proposal coordinates;  $\hat{c}_i$  refers to the confidence of radical prediction, and  $c_i$  is the existence of radical;  $\lambda$  is the weight of absence cases, and we set  $\lambda=0.05$  in the experiments. Then we define  $\mathcal{L}_S$  as the loss function of  $OP_s$  to learn from categories of structural relations:

$$\mathcal{L}_S = \frac{\lambda_s}{N} \sum_{i=0}^N q_i(s) \log(\hat{q}_i(s)), \quad (8)$$

where  $N$  is the number of categories,  $q_i(s)$  is the probability of target structural relation, while  $\hat{q}_i(s)$  is the predicted one. To sum up, the full loss function of RIE can be represented as:  $\mathcal{L}_{RIE} = \mathcal{L}_r + \mathcal{L}_{coo} + \mathcal{L}_{isR} + \mathcal{L}_S$ .

## 4.2 Knowledge Graph Reasoner

Inspired by graph-based recommendation methods [Wang *et al.*, 2020; Shi *et al.*, 2020], we propose KGR, which aims to achieve zero-shot character recognition via a reasoning-based strategy rather than hard-matching strategies. A weight fusion-based reasoning algorithm is proposed to obtain better recognition performance, where a character knowledge graph (CKG) is exploited to organize character information. We also utilize the predictions from RIE as soft labels to improve the adaptive ability of the reasoning process.

**Character Knowledge Graph.** We intend to reuse existing knowledge graphs of characters in various languages, including Oracle, Bronze [Chi *et al.*, 2022], Korean, and simplified Chinese<sup>1</sup>, where we can extract the radical composition of characters and structural relations. To achieve reasoning-based character recognition, we focus on three kinds of entities in these large-scale CKGs, i.e., character, radical, and structure<sup>2</sup>. Note that characters are associated with the corresponding radicals by the relation *contain*, and the structures are connected with characters by *compose*, which provides the possibility of more flexible reasoning via CKG.

**Weight Fusion-based Reasoning.** The inputs of the weight fusion-based reasoning algorithm  $CharReason(\cdot)$  are  $CKG$ , radical predictions ( $RP_s$ ), and structural relation prediction ( $SP$ ). The predictions of radical categories  $RP_s \in \mathbb{R}^{num \times n_r}$  are extracted from the output of the projector  $OP_r$ , where  $num$  is the number of radicals recognized by RIE and  $n_r$  is the length of each prediction.  $SP$  is the output of the projector  $OP_r$  with the length  $n_s$ , which is the number of predefined structural relations.

As shown in the Algorithm 1, firstly, we map the predictions in each  $RP$  to generate the candidate radical mappings  $M$ , where the confidence of each  $m_i \in M$  is calculated by:

$$m_i.conf = \frac{1}{num} \sum_{i=1}^{num} p_r, \quad (9)$$

where  $p_r$  is the prediction confidence *conf* of each candidate radical in the mapping  $m_i$ . We search for candidate characters  $C_r$  that match the radicals in  $m_i$  via  $searchRad(\cdot)$  and candidate characters  $C_s$  that match the relations  $sp_j$  via

---

### Algorithm 1 Weight Fusion-based Reasoning.

$PC = CharReason(CKG, RP_s, SP)$

---

#### Require:

Character knowledge graph  $CKG$ ; Radical predictions  $RP_s$ ;  
Structural relations prediction  $SP$ .

#### Ensure:

Character predictions with confidence  $PC$ .

```

1:  $M = map(RP_s)$ ;
2: for each  $m_i \in maxSort(M)$  do
3:    $C_r = searchRad(m_i, CKG)$ ;
4:   for each  $sp_j \in maxSort(SP)$  do
5:      $C_s = searchStr(sp_j, CKG)$ ;
6:      $C_t = C_r \cap C_s$ 
7:      $p_c = \theta m_i.conf + (1 - \theta) sp_j.conf$ ;
8:      $PC.add(C_t, p_c)$ 
9:   end for
10: end for
11: return  $maxSort(PC)$ 
    
```

---

$searchStr(\cdot)$  in CKG, where  $sp_j \in SP$ . Note that we sort  $M$  and  $SP$  via  $maxSort(\cdot)$  based on the value of prediction confidence before searching to speed up the reasoning process. As a result, we have candidate character  $C_t$  that satisfies both radical mapping and structural relation. Its confidence  $p_c$  is calculated by a weighted fusion of the  $m_i$  and  $sp_j$  in line 7, where  $\theta=0.7$ . The obtained  $C_t$  and the corresponding confidences  $p_c$  are stored in the character prediction  $PC$ . The algorithm outputs the sorted  $PC$  as the final recognition result, where we have the confidence of all candidates  $C_t$  to maximize the possibility of correct recognition.

Our proposed KGR comprehensively considers the extracted character information from RIE, which effectively alleviates the low-precision character reasoning issue caused by hard-matching strategies. Note that KGR supports adding new character categories by updating CKG without additional model training. Thus, we consider RZCR a zero-shot method that is able to recognize unseen character categories.

## 5 Experiments

### 5.1 Experimental Setup

**Datasets.** To evaluate our method on real-world character image sets which suffer from the few-sample problem, we introduce a new character image dataset called OracleRC. We collect oracle rubbing images from [Wu, 2012] and normalize images by a denoising method [Shi *et al.*, 2022a]. OracleRC includes 2,005 character categories that can be decomposed by 202 radicals and 14 structural relations, where the number of character samples ranges from 1 to 32. Radicals and structural relations were manually annotated by 8 linguists. We also validate RZCR on handwritten Chinese datasets IC-DAR2013 [Yin *et al.*, 2013], HWDB1.1 [Liu *et al.*, 2013], scene character dataset CTW [Yuan *et al.*, 2019], and Korean dataset PE92 [KIM *et al.*, 1996] for a comprehensive evaluation. To increase the adaptability of RZCR, we propose a radical splicing-based synthetic character strategy to enlarge the training set and reduce the cost of human annotations.

**Implementation Details.** The resolution of the input image is  $416 \times 416$ . We exploit data enhancement strategies, includ-

<sup>1</sup><http://humanum.arts.cuhk.edu.hk/Lexis/lexi-mf/>

<sup>2</sup>CKGs record the structure of a character, which is considered equal to our defined structural relation.

Method	OracleRC				ICDAR2013		CTW		HWDB1.1	
	Top-1	Top-3	Top-5	Cat <sub>Avg</sub>	Top-1	Cat <sub>Avg</sub>	Top-1	Cat <sub>Avg</sub>	Top-1	Cat <sub>Avg</sub>
AlexNet [Krizhevsky <i>et al.</i> , 2012]	26.93%	36.45%	40.03%	21.74%	89.99%	80.14%	76.49%	61.28%	88.74%	85.32%
VGG16 [Simonyan and Zisserman, 2015]	27.75%	38.12%	41.53%	20.38%	90.68%	82.76%	79.38%	68.34%	89.67%	84.60%
HCCR-GoogLeNet [Zhong <i>et al.</i> , 2015]	28.52%	36.75%	39.86%	18.81%	96.26%	88.97%	82.28%	71.21%	94.85%	90.36%
DropSample-DCNN [Yang <i>et al.</i> , 2016]	29.19%	39.27%	42.03%	19.59%	97.23%	89.10%	82.37%	71.21%	<b>96.57%</b>	91.42%
ResNet [He <i>et al.</i> , 2016]	28.50%	33.02%	40.66%	21.98%	92.18%	85.68%	79.46%	69.82%	90.98%	86.06%
DenseNet [Huang <i>et al.</i> , 2017]	27.85%	38.63%	47.48%	19.20%	95.90%	90.36%	79.88%	68.48%	94.32%	89.72%
DirectMap [Zhang <i>et al.</i> , 2017]	30.48%	<b>44.89%</b>	<b>54.72%</b>	23.59%	<b>97.37%</b>	90.62%	84.23%	72.50%	96.25%	91.28%
M-RBC + IR [Yang <i>et al.</i> , 2017]	30.53%	42.74%	49.32%	20.72%	<b>97.37%</b>	88.70%	83.65%	73.07%	96.14%	90.76%
RAN [Zhang <i>et al.</i> , 2018]	35.37%	-	-	32.48%	93.79%	88.69%	81.80%	76.59%	92.28%	89.43%
DenseRAN [Wang <i>et al.</i> , 2018a]	36.02%	-	-	32.16%	96.66%	<b>91.02%</b>	85.56%	<b>82.47%</b>	95.32%	<b>91.76%</b>
FewshotRAN [Wang <i>et al.</i> , 2019a]	33.31%	-	-	30.90%	96.97%	90.42%	86.78%	81.16%	96.32%	91.59%
HDE-Net [Cao <i>et al.</i> , 2020]	<b>36.79%</b>	-	-	<b>33.10%</b>	96.74%	88.75%	<b>89.25%</b>	78.94%	95.63%	90.48%
Stroke-to-Character [Chen <i>et al.</i> , 2021]	27.30%	-	-	20.09%	96.28%	89.11%	85.29%	77.14%	93.97%	89.83%
STAR [Zeng <i>et al.</i> , 2022]	30.65%	-	-	23.76%	97.11%	88.82%	85.43%	78.26%	94.24%	90.03%
HCRN [Huang <i>et al.</i> , 2022]	35.84%	-	-	32.24%	96.70%	88.97%	85.59%	77.62%	95.86%	91.67%
RZCR (Ours)	<b>61.36%</b>	<b>71.02%</b>	<b>74.39%</b>	<b>58.84%</b>	<b>97.42%</b>	<b>91.43%</b>	<b>88.74%</b>	<b>82.89%</b>	<b>96.45%</b>	<b>92.28%</b>

Table 2: Quantitative comparisons with state-of-the-art methods on four datasets. The best and second-best results are highlighted in red and blue colors, respectively.

Method	OracleRC			Combined dataset			CTW		
	c-500	c-1000	c-1205	c-1000	c-2000	c-2755	c-1000	c-2000	c-3150
DenseRAN [Wang <i>et al.</i> , 2018a]	5.28%	10.67%	11.58%	8.44%	19.51%	30.68%	0.54%	1.95%	5.39%
HDE-Net [Cao <i>et al.</i> , 2020]	7.12%	9.76%	10.51%	12.77%	25.13%	33.49%	2.11%	6.96%	7.75%
Stroke-to-Character [Chen <i>et al.</i> , 2021]	3.37%	7.48%	7.79%	13.85%	25.73%	37.91%	2.54%	6.82%	8.61%
STAR [Zeng <i>et al.</i> , 2022]	6.14%	10.62%	12.23%	19.47%	35.53%	43.86%	3.77%	11.00%	11.27%
RZCR (Ours)	<b>39.21%</b>	<b>52.43%</b>	<b>54.28%</b>	<b>65.58%</b>	<b>73.56%</b>	<b>78.75%</b>	<b>49.73%</b>	<b>61.39%</b>	<b>64.82%</b>

Table 3: Comparisons with zero-shot character recognition methods.

ing translation, rotation, scaling, and background transformation. Parameters  $K=13$  and  $M=3$  are set in RIE. All experiments are conducted with Adadelta optimization where the hyperparameters are set to  $\rho=0.95$  and  $\epsilon=10^{-6}$ .

### 5.2 Experimental Results

We compare RZCR with state-of-the-art character recognition methods on four datasets, and the results are shown in Table 2. We output Top-n prediction by confidence to present the average classification accuracy based on samples. Considering a few categories with more samples are not enough to reflect the overall recognition performance in an unbalanced dataset, we also calculate the average accuracy for each category and then average over all categories, i.e., Cat<sub>Avg</sub>. Note that in Table 2, general OCR methods are presented in the former rows, and the latter are zero-shot recognition methods. For all four datasets, we select 80% of the samples in each character category as the training set and the remaining as the test set. Note that categories containing only one sample are not included in this experiment since general recognition methods are not able to train on these categories.

For the few-sample dataset OracleRC, we can find RZCR significantly outperforms all general OCR methods and zero-shot methods on Top-n accuracy and Cat<sub>Avg</sub>. We also find zero-shot methods except Stroke-to-Character gain higher accuracy than general methods on Top-1 accuracy and Cat<sub>Avg</sub>. Note that lower Cat<sub>Avg</sub> obtained by general OCR methods

means that these methods perform poorly on categories with few samples. The main reason for this disparity is that an insufficient amount of training data limits their performance. In contrast, for zero-shot recognition methods, decomposing characters into elements brings an increasing number of training samples and a decrease in training categories, which alleviates the few-sample issue. Meanwhile, the better performance of RZCR among zero-shot methods benefits from the powerful KGR with a flexible reasoning strategy, as we discussed before. We also evaluate RZCR on ICDAR2013, CTW, and HWDB1.1 to demonstrate its adaptability, where the same training and testing setup are conducted for all methods. As shown in Table 2, both general and zero-shot methods are competitive in these cases with sufficient training samples, where our RZCR also gains promising results. We find RZCR performs better than others on metric Cat<sub>Avg</sub>, which proves our method is less influenced by categories with different numbers of samples during recognition.

Then, we conduct an additional experiment to demonstrate the validity of RZCR on zero-shot character recognition, and the results are presented in Table 3. Only zero-shot recognition methods are included in this experiment since general methods are not able to recognize unseen character categories. Three datasets are applied, including OracleRC, CTW, and a combined dataset by two handwritten character datasets ICDAR2013 and HWDB1.1. We follow the experimental setting introduced by previous zero-shot methods

	Top-1 RPs	Top-1 RPs & SP	RPs	RPs & SP (Ours)
$Acc_C$	38.56%	37.02%	58.06%	<b>61.36%</b>

Table 4: Ablation study on KGR.

	Radicals	Characters	Advances
ICDAR	95.94%	97.42%	+ <b>1.48%</b>
CTW	84.86%	88.74%	+ <b>3.88%</b>
OracleRC	58.59%	61.36%	+ <b>2.47%</b>

Table 5: Comparison of radical and character recognition, to prove the effect of KGR.

	Without FAL	With FAL (Ours)
$Acc_R/Acc_{SR}$	55.87%/70.64%	<b>58.59%/79.41%</b>

Table 6: Comparisons on baseline networks in RIE.

[Wang *et al.*, 2018a; Chen *et al.*, 2021], where we fix the test categories at the beginning and gradually increase the number of training categories. Thus, in this experiment, 800 out of 2,005 categories in OracleRC are unseen categories used for testing. Similarly, 1,000/3,755 and 500/3,650 categories are applied for testing on the combined dataset and CTW, respectively. Meanwhile, as shown in Table 3, *c-m* refers to the comparing group which applied *m* training categories,  $m \in \{500, 1000, \dots\}$ . We can find that our proposed RZCR significantly surpasses other sequence-based zero-shot methods in each set of datasets. More specifically, RZCR benefits from the CKG-based knowledge organization and a flexible reasoning algorithm, which results in the possibility of correct character recognition even with incorrect Top-1 radical prediction. Thus, we can conclude that our proposed RZCR, as a generic character recognition method, is effective for different character datasets, particularly for few-sample datasets.

### 5.3 Ablation Study

**Impact of Reasoning Strategies in KGR.** We also conduct experiments to validate the effectiveness of reasoning-based character recognition. Firstly, we change the radical information and matching strategy in KGR. In Table 4, “Top-1 RPs” refers to a hard-matching strategy that considers only radicals with the highest confidence in RPs, while “Top-1 RPs & SP” utilizes both radicals and structural relations with the highest confidence in RPs and SP, respectively. “RPs” considers all candidate radicals in RPs for reasoning, while “RPs & SP” exploits all candidates in both RPs and SP, i.e., our reasoning-based algorithm. We see that the character recognition accuracy  $Acc_C$  of the latter two is significantly higher than that of the former two, which shows our reasoning-based algorithm surpasses hard-matching strategies. We can also find the increasing number of elements limits the performance of the hard-matching strategy, while the accuracy of our algorithm increases when adding structural relations. A similar phenomenon can also be found in Table 5, where character recognition results by RZCR are higher than radical recognition results only using RIE since the reasoning-based KGR can maximize the possibility of correct recognition.

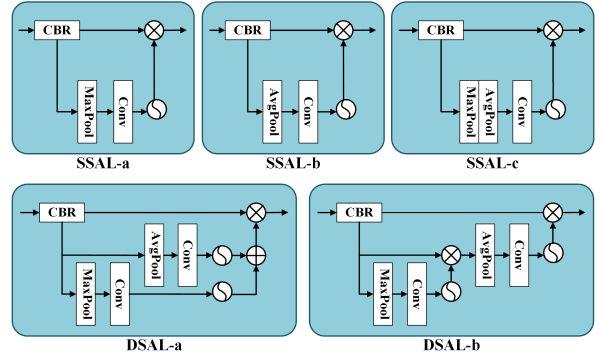


Figure 4: Five connection schemes for the attention layer.

	SSAL-a	SSAL-b	SSAL-c	DSAL-a	DSAL-b (Ours)
$Acc_R$	56.94%	57.18%	57.83%	58.36%	<b>58.59%</b>

Table 7: Ablation study on attention layer schemes.

**Impact of Baseline Networks in RIE.** RP and SP are output by  $OP_r$  and  $OP_s$  via RABs and SRBs, respectively, in RIE, where the two features are interacted by the FAL layer. As shown in Table 6, we apply an ablation test to present the superiority of the proposed FAL. “Without FAL” outputs RPs and SP by two baseline networks, respectively. “With FAL” indicates that RABs and SRBs exchange semantic information by FAL, i.e., RIE.  $Acc_R$  and  $Acc_{SR}$  refer to the recognition accuracy of the radical categories and structural relations. We find that RIE obtains better performance in both metrics. This indicates that the feature extraction of radicals and structures can be improved mutually since the structural relation is relevant to the radical location.

**Impact of Attention Layer Schemes.** In RIE, we aim to utilize the attention mechanism to refine the radical information extraction on character images with overlapping radicals and unclear boundaries. As shown in Fig. 4, we design five possible schemes to stack attention layers, where three schemes are proposed based on the single spatial attention layer (SSAL) and the remaining two are based on the DSAL. The results of the radical recognition by different schemes are shown in Table 7. We conduct this experiment on the OracleRC dataset, where we record the radical recognition accuracy  $Acc_R$ . It is easy to find that DSAL-based schemes outperform three SSAL-based schemes. That is because DSALs consider both foreground information for radical feature selection and background information to maintain the integrity of radicals. Thus, we select better-performed DSAL-b as the attention layer in RABs.

## 6 Conclusion

In this paper, we first introduce the importance of radical information for character recognition and discuss two character decomposition strategies. Then, we propose RZCR, a novel zero-shot character recognition method, to deal with unbalanced character datasets with few samples. RZCR obtains promising experimental results on several datasets, especially on few-sample datasets.

## Acknowledgments

This work is supported by the National Natural Science Foundation of China (No.62077027), the Department of Science and Technology of Jilin Province, China (20230201086GX), the EU’s Horizon 2020 FET proactive project (grant agreement No.823783), and the “Paleography and Chinese Civilization Inheritance and Development Program” Collaborative Innovation Platform (No.G3829). This work is also supported by Professor Chuntao Li and his team from the School of Archeology, Jilin University.

## References

- [Akata *et al.*, 2015] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for image classification. *IEEE TPAMI*, 38(7):1425–1438, 2015.
- [Anderson, 2006] Chris Anderson. *The long tail: Why the future of business is selling less of more*. Hachette Books, 2006.
- [Balaha *et al.*, 2021] Hossam Magdy Balaha, Hesham Arafat Ali, Mohamed Saraya, and Mahmoud Badawy. A new arabic handwritten character recognition deep learning system (ahcr-dls). *Neural Computing and Applications*, 33(11):6325–6367, 2021.
- [Cao *et al.*, 2020] Zhong Cao, Jiang Lu, Sen Cui, and Changshui Zhang. Zero-shot handwritten chinese character recognition with hierarchical decomposition embedding. *Elsevier PR*, 107:107488, 2020.
- [Chen *et al.*, 2021] Jingye Chen, Bin Li, and Xiangyang Xue. Zero-shot chinese character recognition with stroke-level decomposition. In *IJCAI*, 2021.
- [Chi *et al.*, 2022] Yang Chi, Fausto Giunchiglia, Daqian Shi, Xiaolei Diao, Chuntao Li, and Hao Xu. Zinet: Linking chinese characters spanning three thousand years. In *ACL*, pages 3061–3070, 2022.
- [Cireşan and Meier, 2015] Dan Cireşan and Ueli Meier. Multi-column deep neural networks for offline handwritten chinese character classification. In *IJCNN*, 2015.
- [Géron, 2019] Aurélien Géron. *Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow: Concepts, tools, and techniques to build intelligent systems*. O’Reilly Media, 2019.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [Ho *et al.*, 2003] Connie Suk-Han Ho, Ting-Ting Ng, and Wing-Kin Ng. A “radical” approach to reading development in chinese: The role of semantic radicals and phonetic radicals. *Journal of literacy research*, 35(3):849–878, 2003.
- [Huang *et al.*, 2017] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017.
- [Huang *et al.*, 2019] Shuangping Huang, Haobin Wang, Yongge Liu, Xiaosong Shi, and Lianwen Jin. Obc306: A large-scale oracle bone character recognition dataset. In *ICDAR*, 2019.
- [Huang *et al.*, 2022] Guanjie Huang, Xiangyu Luo, Shaowei Wang, Tianlong Gu, and Kaile Su. Hippocampus-heuristic character recognition network for zero-shot learning in chinese character recognition. *Pattern Recognition*, page 108818, 2022.
- [KIM *et al.*, 1996] Dae-Hwan KIM, Young-Sup Hwang, Sang-Tae Park, Eun-Jung Kim, Sang-Hoon Paek, and Sung-Yang BANG. Handwritten korean character image database pe92. *IEICE TOIS*, 79(7):943–950, 1996.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [Lampert *et al.*, 2009] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009.
- [Lampert *et al.*, 2013] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE TPAMI*, 36(3):453–465, 2013.
- [Liu *et al.*, 2011] Cheng-Lin Liu, Fei Yin, Da-Han Wang, and Qiu-Feng Wang. Casia online and offline chinese handwriting databases. In *ICDAR*, 2011.
- [Liu *et al.*, 2013] Cheng-Lin Liu, Fei Yin, Da-Han Wang, and Qiu-Feng Wang. Online and offline handwritten chinese character recognition: benchmarking on new databases. *Elsevier PR*, 46(1):155–162, 2013.
- [Luo *et al.*, 2020] Canjie Luo, Yuanzhi Zhu, Lianwen Jin, and Yongpan Wang. Learn to augment: Joint data augmentation and network optimization for text recognition. In *CVPR*, June 2020.
- [Lyu *et al.*, 2017] Pengyuan Lyu, Xiang Bai, Cong Yao, Zhen Zhu, Tengting Huang, and Wenyu Liu. Auto-encoder guided gan for chinese calligraphy synthesis. In *ICDAR*, 2017.
- [Mushtaq *et al.*, 2021] Faisal Mushtaq, Muzafar Mehraj Misgar, Munish Kumar, and Surinder Singh Khurana. UrduDeepNet: offline handwritten urdu character recognition using deep neural network. *Neural Computing and Applications*, 33(22):15229–15252, 2021.
- [Myers, 2016] James Myers. Knowing chinese character grammar. *Elsevier Cognition*, 147:127–132, 2016.
- [Qu *et al.*, 2018] Xiwen Qu, Weiqiang Wang, Ke Lu, and Jianshe Zhou. Data augmentation and directional feature maps extraction for in-air handwritten chinese character recognition based on convolutional neural network. *Pattern recognition letters*, 111:9–15, 2018.
- [Rohrbach *et al.*, 2011] Marcus Rohrbach, Michael Stark, and Bernt Schiele. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *CVPR*, 2011.



- [Shanthi and Duraiswamy, 2010] N Shanthi and K Duraiswamy. A novel svm-based handwritten tamil character recognition system. *Springer Pattern Analysis and Applications*, 13(2):173–180, 2010.
- [Shen, 2005] Helen H Shen. An investigation of chinese-character learning strategies among non-native speakers of chinese. *Elsevier System*, 33(1):49–68, 2005.
- [Shi et al., 2020] Daqian Shi, Ting Wang, Hao Xing, and Hao Xu. A learning path recommendation model based on a multidimensional knowledge graph framework for e-learning. *Knowledge-Based Systems*, 195:105618, 2020.
- [Shi et al., 2022a] Daqian Shi, Xiaolei Diao, Lida Shi, Hao Tang, Yang Chi, Chuntao Li, and Hao Xu. Charformer: A glyph fusion based attentive framework for high-precision character image denoising. In *ACM MM*, 2022.
- [Shi et al., 2022b] Daqian Shi, Xiaolei Diao, Hao Tang, Xiaomin Li, Hao Xing, and Hao Xu. Rcrn: Real-world character image restoration network via skeleton extraction. In *ACM MM*, 2022.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [Su and Wang, 2003] Yih-Ming Su and Jhing-Fa Wang. A novel stroke extraction method for chinese characters using gabor filters. *Elsevier PR*, 36(3):635–647, 2003.
- [Sun et al., 2021] Xiaohong Sun, Jinan Gu, and Hongying Sun. Research progress of zero-shot learning. *Applied Intelligence*, 51(6):3600–3614, 2021.
- [Wang et al., 2018a] Wenchao Wang, Jianshu Zhang, Jun Du, Zi-Rui Wang, and Yixing Zhu. Denscran for offline handwritten chinese character recognition. In *ICFHR*, 2018.
- [Wang et al., 2018b] Xiaolong Wang, Yufei Ye, and Abhinav Gupta. Zero-shot recognition via semantic embeddings and knowledge graphs. In *CVPR*, 2018.
- [Wang et al., 2019a] Tianwei Wang, Zecheng Xie, Zhe Li, Lianwen Jin, and Xiangli Chen. Radical aggregation network for few-shot offline handwritten chinese character recognition. *Elsevier PRL*, 125:821–827, 2019.
- [Wang et al., 2019b] Wei Wang, Vincent W Zheng, Han Yu, and Chunyan Miao. A survey of zero-shot learning: Settings, methods, and applications. *ACM TIST*, 10(2):1–37, 2019.
- [Wang et al., 2020] Ting Wang, Daqian Shi, Zhaodan Wang, Shuai Xu, and Hao Xu. Mrp2rec: Exploring multiple-step relation path semantics for knowledge graph-based recommendations. *IEEE Access*, 8:134817–134825, 2020.
- [Wu, 2012] Zhenfeng Wu. *Shang and Zhou Bronze Inscriptions and Image Integration*. Shanghai Classics Publishing House, 2012.
- [Xie et al., 2019] Guo-Sen Xie, Li Liu, Xiaobo Jin, Fan Zhu, Zheng Zhang, Jie Qin, Yazhou Yao, and Ling Shao. Attentive region embedding network for zero-shot learning. In *CVPR*, 2019.
- [Yang et al., 2016] Weixin Yang, Lianwen Jin, Dacheng Tao, Zecheng Xie, and Ziyong Feng. Dropsample: A new training method to enhance deep convolutional neural networks for large-scale unconstrained handwritten chinese character recognition. *Pattern Recognition*, 58:190–203, 2016.
- [Yang et al., 2017] Xiao Yang, Dafang He, Zihan Zhou, Daniel Kifer, and C Lee Giles. Improving offline handwritten chinese character recognition by iterative refinement. In *ICDAR*, volume 1, pages 5–10. IEEE, 2017.
- [Yang et al., 2019] Mingkun Yang, Yushuo Guan, Minghui Liao, Xin He, Kaigui Bian, Song Bai, Cong Yao, and Xiang Bai. Symmetry-constrained rectification network for scene text recognition. In *ICCV*, pages 9147–9156, 2019.
- [Yeung et al., 2016] Pui-sze Yeung, Connie Suk-han Ho, David Wai-ock Chan, and Kevin Kien-hoa Chung. Orthographic skills important to chinese literacy development: The role of radical representation and orthographic memory of radicals. *Reading and Writing*, 29(9):1935–1958, 2016.
- [Yin et al., 2013] Fei Yin, Qiu-Feng Wang, Xu-Yao Zhang, and Cheng-Lin Liu. Icdar 2013 chinese handwriting recognition competition. In *ICDAR*, 2013.
- [Yuan et al., 2012] Aiquan Yuan, Gang Bai, Lijing Jiao, and Yajie Liu. Offline handwritten english character recognition based on convolutional neural network. In *DAS*, 2012.
- [Yuan et al., 2019] Tai-Ling Yuan, Zhe Zhu, Kun Xu, Cheng-Jun Li, Tai-Jiang Mu, and Shi-Min Hu. A large chinese text dataset in the wild. *Springer JCST*, 34(3):509–521, 2019.
- [Zatsepin et al., 2019] Michael Zatsepin, Yury Vatlin, Iurii Chulinin, and Aleksei Zhuravlev. Fast korean syllable recognition with letter-based convolutional neural networks. In *ICDARW*, volume 7, pages 10–13. IEEE, 2019.
- [Zeng et al., 2022] Jinshan Zeng, Ruiying Xu, Yu Wu, Hongwei Li, and Jiaying Lu. Star: Zero-shot chinese character recognition with stroke-and radical-level decompositions. *arXiv preprint arXiv:2210.08490*, 2022.
- [Zhan and Lu, 2019] Fangneng Zhan and Shijian Lu. Esir: End-to-end scene text recognition via iterative image rectification. In *CVPR*, pages 2059–2068, 2019.
- [Zhang et al., 2017] Xu-Yao Zhang, Yoshua Bengio, and Cheng-Lin Liu. Online and offline handwritten chinese character recognition: A comprehensive study and new benchmark. *Elsevier PR*, 61:348–360, 2017.
- [Zhang et al., 2018] Jianshu Zhang, Yixing Zhu, Jun Du, and Lirong Dai. Radical analysis network for zero-shot learning in printed chinese character recognition. In *ICME*, 2018.
- [Zhang et al., 2020] Jianshu Zhang, Jun Du, and Lirong Dai. Radical analysis network for learning hierarchies of chinese characters. *Elsevier PR*, 103:107305, 2020.
- [Zhong et al., 2015] Zhuoyao Zhong, Lianwen Jin, and Zecheng Xie. High performance offline handwritten chinese character recognition using googlenet and directional feature maps. In *ICDAR*, pages 846–850. IEEE, 2015.