

# Self-supervised Character-to-Character Distillation for Text Recognition

Tongkun Guan<sup>1</sup>, Wei Shen<sup>1✉</sup>, Xue Yang<sup>1</sup>, Qi Feng<sup>2</sup>, Zekun Jiang<sup>1</sup>, Xiaokang Yang<sup>1</sup>

<sup>1</sup> MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

<sup>2</sup> Department of Automation, Shanghai Jiao Tong University

{gtk0615, wei.shen, yangxue-2019-sjtu, fengqi, zkjiangzekun.cmu, xkyang}@sjtu.edu.cn

## Abstract

When handling complicated text images (e.g., irregular structures, low resolution, heavy occlusion, and uneven illumination), existing supervised text recognition methods are data-hungry. Although these methods employ large-scale synthetic text images to reduce the dependence on annotated real images, the domain gap still limits the recognition performance. Therefore, exploring the robust text feature representations on unlabeled real images by self-supervised learning is a good solution. However, existing self-supervised text recognition methods conduct sequence-to-sequence representation learning by roughly splitting the visual features along the horizontal axis, which limits the flexibility of the augmentations, as large geometric-based augmentations may lead to sequence-to-sequence feature inconsistency. Motivated by this, we propose a novel self-supervised **Character-to-Character Distillation** method, CCD, which enables versatile augmentations to facilitate general text representation learning. Specifically, we delineate the character structures of unlabeled real images by designing a self-supervised character segmentation module. Following this, CCD easily enriches the diversity of local characters while keeping their pairwise alignment under flexible augmentations, using the transformation matrix between two augmented views from images. Experiments demonstrate that CCD achieves state-of-the-art results, with average performance gains of 1.38% in text recognition, 1.7% in text segmentation, 0.24 dB (PSNR) and 0.0321 (SSIM) in text super-resolution. Code is available at <https://github.com/TongkunGuan/CCD>.

## 1. Introduction

Recognizing text from images is a fundamental task in computer vision with applications in various scenarios, such as recognizing latex formulas [41], work-piece serial numbers [18], text logos [38], etc., and contributes significantly to the multi-modal analysis [55] and text understanding [62, 46]. However, existing text recognition meth-

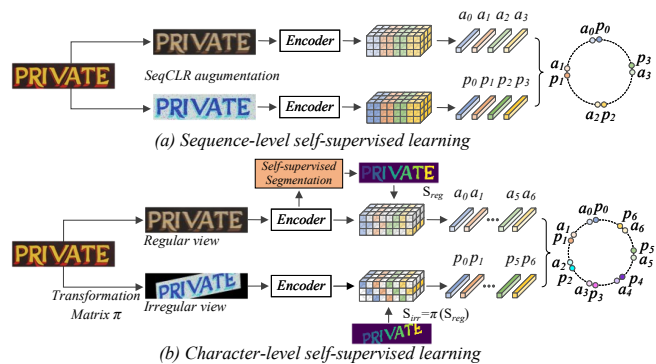


Figure 1. Conceptual illustration of two different self-supervised paradigms. (a) is a sequence-level method that takes feature blocks horizontally-split from a sequence as the basic items for representation learning. (b) is our character-level method that incorporates a self-supervised segmentation head to generate individual character structures ( $S_{reg}$  and  $S_{irr}$ ), and utilizes the resulting character-level features as the basic items for representation learning.

ods [48, 16, 51, 66, 50, 30, 24] are data-hungry, *i.e.*, they require sufficient data with at least text transcriptions for producing accurate character predictions through implicit attention learning [19]. Even more, some supervised attention methods [24, 48, 32] require extra character-level bounding boxes. These annotations on text images are expensive and laborious. Although they employ text synthesis techniques [20, 26] to substitute for labour-intensive text annotation tasks, the domain gap between real and synthetic images still limits the performance of text recognition [61]. Therefore, exploring the potential of unlabeled real text images is of great importance as they are readily available.

Recently, self-supervised learning methods [1, 61, 34] have attracted considerable attention, which attempt to leverage the intrinsic qualities of unlabeled real text images to learn proper visual representations, followed by fine-tuning on text-related downstream tasks with less annotated data. Specifically, SeqCLR [1] adopts the SimCLR [11] framework to ensure sequence-to-sequence consistency between the two augmented views, in which the sequence is composed of several non-overlapping feature blocks horizontally splitting from the visual features of the text im-

age. DiG [61] employs both a sequence-level contrastive learning task [1] and a masked image modeling task [22] to learn feature representations. These methods formulate the learning via sequence-level pretext tasks illustrated in Fig.1 (a). We argue that roughly splitting visual features of text images into a feature sequence along the horizontal axis has two weaknesses: 1) Inflexible data augmentation strategy, as large geometric transformations may cause inconsistency among the corresponding items in the feature sequence generated from different views. However, versatile data augmentations are demanded for self-supervised learning in many previous works [11, 2, 25]; 2) Neglecting character structures, which confuses networks to cause inter-character mixture, further downgrades the perception of semantic clue information in character-centric text images. Thus a suitable self-supervised learning paradigm that is tailored for text images with a diversity of word length is in demand.

To address this issue, we propose a new self-supervised learning paradigm in character-level, named **Character-to-Character Distillation (CCD)**, as shown in Fig.2, which enables feature representation consistency across various augmentations by organizing text images into entities, *i.e.*, each character and background regions. Specifically, two views are first generated from each input image: a regular view with color jitter and an irregular view with additional geometric transformations. Each view is fed into the encoder of the student-teacher branches for extracting features that represent the whole view. Then, character regions from regular view are delineated by a joint self-supervised text segmentation and density-based spatial clustering task, and those from irregular view are generated using the known transformation matrix between two views. In this way, CCD naturally ensures the consistency of corresponding characters across views and branches. Consequently, by enjoying the pairwise diversity of local characters under flexible augmentations, CCD effectively enhances the robustness and generalization of the learned features, making it more suitable for text-related downstream tasks. In summary, the main contributions are as follows:

- We propose a novel self-supervised method customized for character-centric text images with a diversity of word length, termed CCD. Different from prior works with sequence-to-sequence pretext tasks, CCD delineates the character structures to establish character-to-character feature representation consistency, enjoying significant augmentation flexibility for extracting general text feature representations.
- CCD shows its prominent superiority in self-supervised representation learning, and consistently and significantly outperforms the state-of-the-art DiG [61] by an average of 1.38% in text recognition, 1.7% in text segmentation, 0.24 dB (PSNR) and

0.0321 (SSIM) in text image super-resolution task, with the same parameters and latency.

## 2. Related Work

### 2.1. Text Recognition

Given a text image and supervised by its text annotations, text recognition methods aim to predict these characters. Specifically, these methods can be roughly summarized into language-free and language-aware methods. For language-free methods [45, 43, 31, 30, 68, 48, 32, 37, 5, 6, 56], they view text recognition as a character classification task and focus on how to extract robust visual features of characters. For example, some works [45, 43, 31, 30, 33, 68, 60] develop implicit attention mechanisms by computing the similarity of feature patches to extract the important items of visual features, *w.r.t.* the current decoding character. Some works [48, 32, 37] employ extra character-level bounding box annotations as the spatial locations of characters to supervise the corresponding attention at the decoding stage, which alleviates the alignment-drifted problem [64] and improves attention correctness. For language-aware methods [16, 8, 53, 63], they execute semantic reasoning on texts with linguistic context to explore the semantic relationships among character, sub-word and word. For example, Fang *et al.* [16] propose a language model by predicting the masked character in a text with linguistic context. Li *et al.* [53] mask the visual features of some characters based on character attention to predict the corresponding character categories.

### 2.2. Self-supervised Learning for Text Recognition

Recently, self-supervised learning [9, 11, 69, 22, 23, 57] objectives have gained considerable traction due to their powerful feature representations for transferring downstream tasks. These methods leverage the intrinsic qualities of unlabeled real images to learn general feature representations, by ensuring the feature consistency of various augmented views. For example, two popular self-supervised pretext tasks in computer vision are designed to establish representation learning: discriminative task [1, 23] and generative task [22, 57].

Inspired by these methods, some self-supervised text recognition methods [1, 61, 34, 65] are designed for text images with indefinite-length characters, which differs from well-curated images with an atomic input element. Specifically, SeqCLR [1] is the first to apply a self-supervised method to text recognition, which proposes a sequence-to-sequence contrastive learning framework on text images. It horizontally splits the visual features into a sequence with fixed-length feature blocks, and each item of the feature sequence from two augmented views is aligned. Following SeqCLR with sequence-level representative learning, DiG [61] updates one of the augmented views as a masked




Figure 2. Overview of self-supervised Character-to-Character Distillation (CCD). Both two augmented views ( $\mathbf{X}_{reg}$  and  $\mathbf{X}_{irr}$ ) are fed into the student branch and the teacher branch to obtain character features, which are represented by  $\mathbf{R}_s$  and  $\mathbf{I}_s$  in the student branch, and  $\mathbf{R}_t$  and  $\mathbf{I}_t$  in the teacher branch. We then build character-to-character representation consistency across different views and different branches.

view and adds a masked image modeling task. PerSec [34] performs hierarchical contrastive learning across each element of features for text recognition. Our work differs from prior works in that we delineate character structures and propose a character-to-character distillation task to learn more universal text features in the representative space.

### 3. Methodology

#### 3.1. Architecture

**Data Augmentation:** The input image  $\mathbf{X}$  undergoes color-based augmentations (e.g., color jitter, color dropout, and grayscale conversion) to create a regular view  $\mathbf{X}_{reg}$ , and a combination of both color- and geometry-based augmentations (e.g., affine transformation and perspective warping) to generate an irregular view  $\mathbf{X}_{irr}$ .

**Encoder  $\mathcal{F}(\cdot)$ :** ViT [14] is employed as the encoder of our method CCD due to its prominent superiority in extracting visual features. Specifically, two augmented views ( $\mathbf{X}_{reg}$  and  $\mathbf{X}_{irr}$ ) are split into non-overlapping patches with size  $4 \times 4$ , and then fed into multi-layer transformer blocks to extract text features.

**Self-supervised Character Segmentation Head  $\Phi(\cdot)$ :** We employ an Unet-like network structure to implement pixel-level text segmentation, which assigns each pixel a foreground or background label. Then, we delineate the character structures from the resulting text segmentation map by a density-based spatial clustering task.

**Patch Head  $\mathcal{H}(\cdot)$ :** takes the character structure regions and text features as inputs and generates character-level feature representations by a mean-pooling operation.

**Projection Head  $\mathcal{P}(\cdot)$ :** consists of a three-layer MLP and a weight-normalized fully connected layer [9], producing the final character features.

**Student & Teacher Branch** contain all of the above units, except for self-supervised character segmentation head  $\Phi(\cdot)$ , which is exclusively designed in the student branch to generate the character segmentation results ( $\mathbf{S}_{reg}$

---

#### Algorithm 1: Pytorch pseudo-code of CCD.

---

```

#  $\mathcal{G}_s, \mathcal{G}_t$ : student and teacher branches
#  $\mathcal{F}_s, \Phi_s$ : encoder and self-supervised character
# segmentation head in the student branch
 $\mathcal{G}_t.params = \mathcal{G}_s.params$ 
for X in loader: # load a minibatch X
# augmentation: regular view and irregular view
 $\mathbf{X}_{reg}, \mathbf{X}_{irr} = \text{augment\_reg}(\mathbf{X}), \text{augment\_irr}(\mathbf{X})$ 
# character regions in the student branch
 $\mathbf{S}_{reg} = \Phi_s(\mathcal{F}_s(\mathbf{X}_{reg}))$ ,  $\mathbf{S}_{irr} = \pi_{irr}(\mathbf{S}_{reg})$ 
# character regions in the teacher branch
 $\mathbf{T}_{seg} = \mathbf{S}_{reg}$ ,  $\mathbf{T}_{irr} = \mathbf{S}_{irr}$ 
# character features in the student branch
 $\mathbf{R}_s, \mathbf{I}_s = \mathcal{G}_s(\mathbf{X}_{reg}, \mathbf{S}_{reg}), \mathcal{G}_s(\mathbf{X}_{irr}, \mathbf{S}_{irr})$ 
# character features in the teacher branch
 $\mathbf{R}_t, \mathbf{I}_t = \mathcal{G}_t(\mathbf{X}_{reg}, \mathbf{T}_{seg}), \mathcal{G}_t(\mathbf{X}_{irr}, \mathbf{T}_{irr})$ 
loss =  $\xi(\mathbf{I}_t, \mathbf{R}_s) + \xi(\mathbf{R}_t, \mathbf{I}_s)$ 
loss.backward() # back-propagate
# student and teacher updates
update( $\mathcal{G}_s$ ) # AdamW optimizer
 $\mathcal{G}_t.params = \lambda * \mathcal{G}_t.params + (1 - \lambda) * \mathcal{G}_s.params$ 

```

---

in regular view and  $\mathbf{S}_{irr}$  in irregular view). For simplicity, the segmentation head of the teacher branch is canceled by utilizing the segmentation results ( $\mathbf{S}_{reg}$  and  $\mathbf{S}_{irr}$ ) as the corresponding character regions ( $\mathbf{T}_{reg}$  in regular view and  $\mathbf{T}_{irr}$  in irregular view). Subsequently, character features ( $\mathbf{R}_s$  and  $\mathbf{I}_s$ ) are generated in the student branch to match the character features ( $\mathbf{R}_t$  and  $\mathbf{I}_t$ ) distribution from the teacher branch. The whole pipeline is illustrated in Fig. 2 and the detailed network is described in Sec. 3.2. Additionally, we also provide a pseudo-code implementation of our character-to-character self-supervised learning method to further illustrate our pipeline as shown in Algo. 1.

#### 3.2. Character-level Representation Learning

In this section, we will illustrate how to obtain character structures and ensure character-to-character consistency for representation learning, which is different from existing sequence-to-sequence self-supervised methods.

**1) Self-supervised Character Segmentation.** Given an unlabeled text image, our objective is to perform instance-level character segmentation, which identifies all character regions and produces a mask for each of them. Specifically,



Figure 3. The self-supervised character segmentation pipeline.

to make the task more feasible and reasonable, we divide it into two sub-tasks: a self-supervised text segmentation task and a clustering-based character segmentation task.

For the self-supervised text segmentation task, we first compute the pseudo-labels  $\mathbf{M}_{pl}$  of the input images and then use it to train our text segmentation network, which assigns a foreground or background category to each pixel. To achieve this, we select a simple and effective  $K$ -means algorithm [21] (setting  $K = 2$ ) to cluster the pixels of each image into a text region (center) and a background region (surrounding) according to the gray values of pixels. Subsequently, the segmentation network employs the outputs of the 2-nd, 4-th, and 6-th layer of the encoder ViT as  $\mathbf{P}_0, \mathbf{P}_1, \mathbf{P}_2$ , and the implementation detail is as follows:

$$\begin{cases} \mathbf{O}_i = \varphi(\mathbf{P}_i), i = 0, 1, 2, \\ \mathbf{O} = \mathcal{T}([\mathbf{O}_0, \mathbf{O}_1, \mathbf{O}_2]), \end{cases} \quad (1)$$

where  $\varphi(\cdot)$  denotes two convolutional layers with Batch-Norm and ReLU activation functions.  $\mathcal{T}(\cdot)$  refers to two  $2 \times$  upsampling operations to restore the resolution of the input image.  $[\cdot]$  represents the concatenation operation along the channel axis. Finally, the network predictions for text segmentation map  $\mathbf{M}_{seg}$  are generated by applying a convolutional layer on  $\mathbf{O}$  for binary classification. A cross-entropy loss  $\mathcal{L}_{seg}$  between  $\mathbf{M}_{pl}$  and  $\mathbf{M}_{seg}$  is employed to optimize our self-supervised text segmentation network.

For now, let us assume that the text segmentation map  $\mathbf{M}_{seg}$  is obtained, the clustering-based character segmentation task aims to get a mask for each of its characters. One can observe that individual characters retain internal pixel-level connectivity, while the spaces between characters exhibit discontinuities in most natural scene text images. Taking advantage of the observation, we employ a density-based spatial clustering method [15] to segment the  $\mathbf{M}_{seg}$  into several clusters. Specifically, a cluster is formed by aggregating all nearby points connected by density and grouping them together. Then, the points within each cluster can be viewed as a character structure. A discussion about hyper-parameters is deferred to ablations and analysis in Sec. 5.

As shown in Fig. 3, the self-supervised character segmentation head exclusively tasks regular views in the student branch as inputs to generate final character segmentation results  $\mathbf{S}_{reg} = [s_{r1}, s_{r2}, \dots, s_{rl}]$  for simplicity in the experiment.  $l$  refers to the number of cluster centers, as well as the word length of the image ideally.

**2) Corresponding Regions Alignment.** An effective data augmentation strategy is crucial to implementing represen-

tation learning [11, 2, 25]. However, in the sequence-level self-supervised text recognition methods, strong geometric transformations will result in the corresponding item inconsistency of feature sequences between different views.

Motivated by this, we propose an alignment strategy for character regions under flexible augmentations. Specifically, for the regular view in the student branch, we have obtained character segmentation results  $\mathbf{S}_{reg}$  as above. To execute character region alignment, other character segmentation results ( $\mathbf{S}_{irr}$ ,  $\mathbf{T}_{reg}$ , and  $\mathbf{T}_{irr}$ ) are generated as follows.

To calculate  $\mathbf{S}_{irr}$  from the irregular view in the student branch, defined the transformation matrix of augmentation as  $\pi$ , the problem can be formulated as: given  $\mathbf{X}_{reg} = \pi_{reg}(\mathbf{X})$ ,  $\mathbf{X}_{irr} = \pi_{irr}(\mathbf{X})$ , and  $\mathbf{S}_{reg}$ , compute  $\mathbf{S}_{irr} = \pi_{reg \rightarrow irr}(\mathbf{S}_{reg})$ . Consequently,  $\mathbf{S}_{irr} = \pi_{irr}(\mathbf{S}_{reg})$  due to  $\pi_{reg \rightarrow irr} = \pi_{irr}(\pi_{reg}^{-1}) = \pi_{irr}$ , where  $\pi_{reg}$  is an identity matrix as the regular view undergoes color-based augmentation pipeline.

For obtaining  $\mathbf{T}_{reg}$  from regular views and  $\mathbf{T}_{irr}$  from irregular views in the teacher branch respectively, the character segmentation results in the student branch are directly employed, *i.e.*,  $\mathbf{T}_{reg} = \mathbf{S}_{reg}$ ,  $\mathbf{T}_{irr} = \mathbf{S}_{irr}$ , due to the same augmentations used in both student and teacher branches.

In this way, our method naturally ensures the alignment of corresponding character regions across different views and branches, which enriches their pairwise diversity. Thus by addressing the challenge of learning feature consistency across diverse augmentations, our method effectively enhances the robustness and generalizability of the learned features, making them more suitable for downstream tasks.

**3) Character-to-character Distillation.** With the aforementioned foundation, we can proceed to implement the character-to-character distillation across different views ( $\mathbf{X}_{reg}$  and  $\mathbf{X}_{irr}$ ) and branches (student and teacher).

Specifically, we first calculate character feature representations ( $\mathbf{R}_s$ ,  $\mathbf{I}_s$ ,  $\mathbf{R}_t$ , and  $\mathbf{I}_t$ ). Taking the regular view  $\mathbf{X}_{reg}$  in the student branch as an example, we obtain the encoded features  $\mathbf{h}_{reg} = \mathcal{F}(\mathbf{X}_{reg})$  and character segmentation results  $\mathbf{S}_{reg} = \Phi(\mathbf{h}_{reg}) = [s_{r1}, \dots, s_{rl}]$ , the patch head  $\mathcal{H}(\cdot)$  further yields the character-level features  $\mathbf{V}_{reg} = [\mathbf{v}_{r1}, \dots, \mathbf{v}_{rl}]$  by a mean-pooling operation as follows:

$$\mathbf{v}_{ri} = \frac{1}{\sum_{x,y} s_{ri}^{(x,y)}} \sum_{x,y} s_{ri}^{(x,y)} \mathbf{h}_{reg}^{(x,y)}, \quad (2)$$

where  $i$  represents the channel index, and  $(x, y)$  indicates the coordinate point in the map.  $\mathbf{V}_{reg}$  is then fed into our projection head  $\mathcal{P}(\cdot)$  to get the final character features  $\mathbf{R}_s = \mathcal{P}(\mathbf{V}_{reg})$  where  $\mathbf{R}_s \in \mathbb{R}^{l \times n}$ . Specifically,  $\mathcal{P}(\cdot)$  includes four linear layers, with GELU layer applied to the first two linear layers having a hidden dimension of 2048, and a normalization operation [9] applied to the third linear layer having an output dimension of 256. The last linear layer projects the features into a high-dimension space with  $n$  dimensions ( $n = 65536$ ).

Table 2. The encoder and decoder configurations. ViT-series refers to three variants of ViT (Tiny, Small, and Base).

Pre-training Stage	Encoder	Encoder Configuration		
		Embed_dim	Depth	Heads
CCD	ViT-Tiny	192	12	3
	ViT-Small	384	12	6
	ViT-Base	512	12	8

---

Fine-tuning Stage	Encoder	Decoder Configuration		
		Embed_dim	Depth	Heads
Text Recognition	ViT-series	512	6	8
Text Segmentation	ViT-Small	384	3	2
Text Image Super-Resolution	ViT-Small	384	3	2

Following the same principle, we can obtain the remaining character features, *i.e.*,  $\mathbf{I}_s \in \mathbb{R}^{l \times n}$  from the irregular view in the student branch,  $\mathbf{R}_t, \mathbf{I}_t \in \mathbb{R}^{l \times n}$  from the regular and irregular view in the teacher branch, respectively.

Inspired by DINO [9], the student is then distilled from the teacher, by optimizing the character-level features ( $\mathbf{R}_s, \mathbf{I}_s$ ) of the student branch to match those ( $\mathbf{R}_t, \mathbf{I}_t$ ) of the teacher branch. Assuming  $\mathbf{a}, \mathbf{b} \in \mathbb{R}^{l \times n}$ , let us define:

$$\xi(\mathbf{a}, \mathbf{b}) = - \sum_{i=1}^l \sum_{j=1}^n \left( \frac{\exp(a_i^j / \tau)}{\sum_{k=1}^n \exp(a_i^k / \tau)} \log \frac{\exp(b_i^j / \tau)}{\sum_{k=1}^n \exp(b_i^k / \tau)} \right)$$

The distillation loss is formulated as:  $\mathcal{L}_{\text{dis}} = \xi(\mathbf{R}_t, \mathbf{I}_s) + \xi(\mathbf{I}_t, \mathbf{R}_s)$ , where  $\tau$ , a temperature parameter, represents  $\tau_s$  and  $\tau_t$  in the student and teacher branches respectively. Finally, the teacher weights  $\theta_t$  are updated by applying an exponential moving average (EMA) on the student weights  $\theta_s$ , which is summarized as  $\theta_t = \lambda \theta_t + (1 - \lambda) \theta_s$ .

### 3.3. Down-stream Tasks

Inherited from the encoder of CCD, the same decoder structures with self-supervised methods [61] for different downstream tasks are employed for fair comparisons. The detailed network configurations are shown in Table 2.

**Text Recognition** adds a transformer-based decoder [30], which consists of 6 transformer blocks and a linear prediction layer with 96 channels to predict characters.

**Text Segmentation** introduces a decoder with 3 transformer blocks and a linear prediction layer. The final dimension is 2 (foreground and background categories).

**Text Image Super-Resolution** The same decoder structure with the text segmentation network is employed, except that the final prediction dimension is replaced by 3 to recover the input image with RGB channels.

## 4. Experiment

### 4.1. Dataset

**Unlabeled Real Data (URD)** contains 15.77M real-world text images, which are cropped from the large-scale Conceptual Captions Dataset<sup>1</sup> by applying the text bounding box results provided by Microsoft Azure OCR system.

**Synthetic Text Data (STD)** consists of two large-scale SynthText [20] (8M) and Synth90k [26] (9M).

<sup>1</sup><https://github.com/google-research-datasets/conceptual-captions>

**Annotated Real Data (ARD)** is collected from natural scenes and contains 2.78M text images (0.71M in TextOCR [47] and 2.07M in Open Image Dataset v5<sup>2</sup>).

**Scene Text Recognition Benchmarks** include three regular text datasets (*i.e.*, IIIT5K-Words (IIIT) [39], ICDAR2013 (IC13) [28], and Street View Text (SVT) [49]) and three irregular text datasets (*i.e.*, ICDAR2015 (IC15) [27], SVT Perspective (SVTP) [42], and CUTE80 (CT) [44]). IIIT, SVT, IC13, IC15, SVTP, and CT benchmarks contain 3000, 647, 1015, 1811, 645, and 288 images, respectively.

**Text Segmentation Benchmark** TextSeg [58] provides 4024 fine-annotated text images. We crop these images and segmentation maps to construct a text instance segmentation dataset according to its word-level bounding box annotations (training set: 10226, testing set: 3445).

**Text Image Super-Resolution Benchmark** TextZoom [52] consists of image pairs with high-resolution and low-resolution text images. Specifically, 17367 image pairs are used for training, while 1619, 1411, and 1343 image pairs are used for evaluation according to difficulty, respectively.

### 4.2. Implementation Details

**Self-supervised Pre-Training** The pre-training experiments are conducted on URD and STD without labels, at a resolution of  $32 \times 128$ , for fair comparisons. Specifically, we employ ViT-series (*i.e.*, ViT-Tiny, ViT-Small, and ViT-Base) as the baseline structure of CCD. We train our CCD with the AdamW optimizer [35], a cosine learning rate scheduler [17] with a base learning rate of  $5e-4$ , a cosine weight delay scheduler [17] from 0.04 to 0.4, batch size with 288, and warm-up for 0.3 epoch in a total of 3 epochs. The temperature  $\tau_s$  and  $\tau_t$  are set to 0.1 and 0.04, respectively. The coefficient  $\lambda$  follows a cosine scheduler [17] from 0.996 to 1.

**Text Recognition Fine-Tuning** Our text recognition network is fine-tuned at  $32 \times 128$  resolution with STD or ARD, and the total training epochs are 10 or 35. The batch size is 384 and the warm-up time is 1 epoch. The same optimizer and learning scheduler are employed.

**Text Segmentation Fine-Tuning** The text segmentation task is fine-tuned at  $32 \times 128$  resolution with the TextSeg dataset. The batch size is 384, the total number of fine-tuning epochs is 800, and the warm-up time is 50 epochs. The same optimizer and learning scheduler are employed.

**Text Image Super-Resolution Fine-Tuning** The batch size is 384, the total number of fine-tuning epochs is 300, and the warm-up time is 100 epochs. The general Peak Signal-to-Noise Ratio (PSNR) and Similarity Index Measure (SSIM) [54] are employed to evaluate the quality of super-resolution images. All experiments are implemented on a server with 3 NVIDIA 3090 GPUs in PyTorch.

<sup>2</sup>[https://storage.googleapis.com/repositories/opencvino\\_training\\_extensions/datasets/open\\_images\\_v5\\_text](https://storage.googleapis.com/repositories/opencvino_training_extensions/datasets/open_images_v5_text)

Table 3. Text recognition results compared to other self-supervised text recognizers. \* means using an extra 100M images for pre-training.

Method	Venue	Data	IIIT	SVT	IC13	IC15	SVTP	CT	Avg.	Params.
SeqCLR [1]	CVPR'21	STD	82.9	-	87.9	-	-	-	-	-
SimAN [36]	CVPR'22	STD	87.5	-	89.9	-	-	-	-	-
PerSec-ViT* [34]	AAAI'22	STD	88.1	86.8	94.2	73.6	77.7	72.7	83.77	-
DiG-ViT-Tiny [61]	MM'22	STD	95.8	92.9	96.4	84.8	87.4	86.1	91.83	20M
CCD-ViT-Tiny	-	STD	96.5(+0.7)	93.4(+0.5)	96.3(-0.1)	85.2(+0.4)	89.8(+2.4)	89.2(+2.9)	92.57(+0.74)	20M
DiG-ViT-Small [61]	MM'22	STD	96.7	93.4	97.1	87.1	90.1	88.5	93.23	36M
CCD-ViT-Small	-	STD	96.8(+0.1)	94.4(+1.0)	96.6(-0.5)	87.3(+0.2)	91.3(+1.2)	92.4(+3.9)	93.59(+0.37)	36M
DiG-ViT-Base [61]	MM'22	STD	96.7	94.6	96.9	87.1	91.0	91.3	93.49	52M
CCD-ViT-Base	-	STD	97.2(+0.5)	94.4(-0.2)	97.0(+0.1)	87.6(+0.5)	91.8(+0.8)	93.3(+2.0)	93.96(+0.47)	52M
DiG-ViT-Tiny [61]	MM'22	ARD	96.4	94.4	96.2	87.4	90.2	94.1	93.37	20M
CCD-ViT-Tiny	-	ARD	97.1(+0.7)	96.0(+1.6)	97.5(+1.3)	87.5(+0.1)	91.6(+1.4)	95.8(+1.7)	94.18(+0.81)	20M
DiG-ViT-Small [61]	MM'22	ARD	97.7	96.1	97.3	88.6	91.6	96.2	94.69	36M
CCD-ViT-Small	-	ARD	98.0(+0.3)	96.4(+0.3)	98.3(+1.0)	90.3(+1.7)	92.7(+1.1)	98.3(+2.1)	95.57(+0.88)	36M
DiG-ViT-Base [61]	MM'22	ARD	97.6	96.5	97.6	88.9	92.9	96.5	94.92	52M
CCD-ViT-Base	-	ARD	98.0(+0.4)	97.8(+1.3)	98.3(+0.7)	91.6(+2.7)	96.1(+3.2)	98.3(+1.8)	96.30(+1.38)	52M

Table 4. Comparison results of scene text recognition methods. “V” and “L” refer to the language-free and language-aware methods, respectively. The best results are shown in bold font. “Avg1” denote the weighted average results of IIIT, SVT, IC13, SVTP and CT by size. “Avg2” denote the weighted average results of IIIT, SVT, IC15, SVTP and CT by size.

Methods	Type	Venue	Data	IIIT	SVT	IC13	IC15	SVTP	CT	Avg1	Avg2	Params.	Time (ms)
PIMNet [43]	V	MM'21	STD	95.2	91.2	93.4	83.5	84.3	84.4	92.60	89.89	-	28.4
TRBA [3]		CVPR'21	STD	92.1	88.9	93.1	78.3	79.5	78.2	89.74	85.97	50M	27.6
PREN2D [59]		CVPR'21	STD	95.6	94.0	-	83.0	87.6	91.7	-	90.88	-	67.4
Text is Text [7]		ICCV'21	STD	92.3	89.9	-	-	84.4	86.3	-	-	-	-
SGBANet [68]		ECCV'22	STD	95.4	89.1	95.1	-	83.1	88.2	92.83	-	-	-
CornerTransformer [?]		ECCV'22	STD	95.9	94.6	96.4	-	91.5	92.0	95.13	-	86M	294.9
MGP-STR [50]		ECCV'22	STD	96.4	94.7	-	87.2	91.0	90.3	-	92.80	148M	12.3
SIGA [19]		CVPR'23	STD	96.6	95.1	96.8	86.6	90.5	93.1	95.58	92.84	113M	56.3
SRN [63]		CVPR'20	STD	94.8	91.5	95.5	82.7	85.1	87.8	93.07	89.74	49M	26.9
ABINet [16]	CVPR'21	STD+WiKi	96.2	93.5	-	86.0	89.3	89.2	-	92.02	-	37M	33.9
JVSR [8]	ICCV'21	STD	95.2	92.2	95.5	-	85.7	89.7	93.53	-	44M	26.3	
VisionLAN [53]	ICCV'21	STD	95.8	91.7	95.7	83.7	86.0	88.5	93.80	90.64	33M	-	
S-GTR [24]	AAAI'22	STD+WiKi	95.8	94.1	-	84.6	87.9	92.3	-	91.50	42M	18.8	
ABINet+ConCLR [66]	AAAI'22	STD+WiKi	96.5	94.3	-	85.4	89.3	91.3	-	92.17	-	-	
PARSeq [4]	ECCV'22	STD	97.0	93.6	96.2	86.5	88.9	92.2	95.28	92.65	-	-	
LevOCR [12]	ECCV'22	STD	96.6	92.9	-	86.4	88.1	91.7	-	92.26	109M	119.0	
CCD-ViT-Tiny	V	-	STD	96.5	93.4	96.3	85.2	89.8	89.2	94.96	91.98	20M	43.2
CCD-ViT-Small		-	STD	96.8	94.4	96.6	87.3	91.3	92.4	95.63	93.11	36M	44.2
CCD-ViT-Base		-	STD	97.2	94.4	97.0	87.6	91.8	93.3	<b>96.02</b>	<b>93.48</b>	52M	45.0
CCD-ViT-Tiny	V	-	ARD	97.1	96.0	97.5	87.5	91.6	95.8	96.34	93.65	20M	43.2
CCD-ViT-Small		-	ARD	<b>98.0</b>	96.4	<b>98.3</b>	90.3	92.7	<b>98.3</b>	97.27	95.13	36M	44.2
CCD-ViT-Base		-	ARD	<b>98.0</b>	<b>97.8</b>	<b>98.3</b>	<b>91.6</b>	<b>96.1</b>	<b>98.3</b>	<b>97.83</b>	<b>95.99</b>	52M	45.0

### 4.3. Experiment Results

**Self-supervised text recognition.** In Table 3, we evaluate the robustness of feature representations for text recognition by comparing the proposed CCD-ViT-series (*i.e.*, Tiny, Small, Base) with previous self-supervised text recognition methods. Our method achieves a new state-of-the-art performance on both regular and irregular benchmarks.

Specifically, our CCD-ViT-Tiny outperforms the sequence-to-sequence method SeqCLR by 13.6% and 8.4% on IIIT and IC13 benchmarks, respectively. Compared to the “PerSec-ViT\*” method that exploits 100M private unlabeled real images for pre-training, CCD-ViT-series is pre-trained on the URD benchmark with 15.77M unlabeled images, and still yields 8.80%, 9.82%, and 10.19% on average accuracy, respectively.

We further conduct a comparison with the previous state-

of-the-art self-supervised method, DiG, using the same pre-training data and network parameters. Specifically, when fine-tuning with the STD, CCD-ViT-series get better text recognition performances than DiG-ViT-series by 0.74%, 0.37%, and 0.47% on average accuracy, respectively. When fine-tuning with the ARD, CCD-ViT-series consistently and significantly achieves performance gains of 0.81%, 0.88%, and 1.38% on average accuracy compared to DiG-ViT-series. These results demonstrate that our proposed character-level representation learning paradigm is superior to the existing sequence-to-sequence self-learning paradigm, particularly on real-world datasets.

**Scene text recognition.** In Table 4, we present a comparison between CCD and previous state-of-the-art (SOTA) supervised text recognition methods. Specifically, CCD-ViT-Tiny achieves competitive recognition results (94.96% vs. 95.58%) with a minimum parameter of 20M. CCD-ViT-

Table 5. The super-resolution evaluation results on the TextZOOM benchmark.

Method	Loss	SSIM				PSNR			
		Easy	Medium	Hard	Avg.	Easy	Medium	Hard	Avg.
Bicubic	$\times$	0.7884	0.6254	0.6592	0.6961	22.35	18.98	19.39	20.35
SRCNN [13]	$L_2$	0.8379	0.6323	0.6791	0.7227	23.48	19.06	19.34	20.78
SRResNet [29]	$L_2+L_{tv}+L_p$	0.8681	0.6406	0.6911	0.7403	24.36	18.88	19.29	21.03
HAN [40]	$L_2$	0.8691	0.6537	0.7387	0.7596	23.30	19.02	20.16	20.95
TSRN [52]	$L_2+L_{Gp}$	0.8897	0.6676	0.7302	0.7690	25.07	18.86	19.71	21.42
TBSRN [10]	$L_{POS}+L_{CON}$	0.8729	0.6455	0.7452	0.7603	23.46	19.17	19.68	20.91
PCAN [67]	$L_2+L_{EG}$	0.8830	0.6781	0.7475	0.7752	24.57	19.14	20.26	21.49
Scratch-ViT-Small	$L_2$	0.8143	0.6288	0.6845	0.7156	22.90	19.65	20.45	21.10
DiG-ViT-Small [61]	$L_2$	0.8613	0.6561	0.7215	0.7522	23.98	19.85	20.57	21.60
CCD-ViT-Small (ours)	$L_2$	0.8822	0.7005	0.7543	<b>0.7843</b>	24.40	20.12	20.18	<b>21.84</b>

Table 6. The text segmentation results on the TextSeg benchmark.

Method	Scratch-ViT-Small	DiG-ViT-Small [61]	CCD-ViT-Small
IoU(%)	78.1	83.1	<b>84.8</b>

Table 7. Ablation study on the text segmentation experiments.

Method	$K$ -means		$M_{pl}$	Self-supervised
	$\Theta$	$1 - \Theta$		
IoU(%)	39.8	40.1	<b>70.0</b>	<b>73.6</b>

Small gets higher performance and outperforms previous SOTA by 0.27% on average accuracy while exhibiting a cost-quality trade-off between accuracy and general evaluation metrics (parameter count and latency). CCD-ViT-Base further refreshes the best text recognition results, achieving average gains of 0.44% and 0.64% with a smaller model size (52M vs. 113M). When fine-tuning on annotated real data, CCD-ViT-Base attains significantly higher accuracy and establishes new SOTA results, with performance gains by 1.0%, 2.7%, 1.5%, 4.4%, 4.6%, and 5.2% compared to the best results from previous methods on IIIT, SVT, IC13, IC15, SVTP, and CT benchmarks. These results underscore the potential of our proposed self-supervised character-level learning as a powerful and flexible method for text recognition.

**Text image super-resolution.** In Table 5, CCD is also applied to text image super-resolution task. Compared to the Scratch-ViT-Small without the pre-training stage, our CCD-ViT-Small gets better super-resolution results on the SSIM and PSNR metrics. Compared to the self-supervised representation learning method DiG, our method consistently leads performance improvements on the SSIM (0.7843 vs. 0.7522) and PSNR (21.84 vs. 21.60) metrics, while using the same parameters and data for pre-training and fine-tuning. Notably, our method also achieves higher performance than previous state-of-the-art super-resolution methods, despite employing just three transformer units connected to the ViT-Small structure without additional designs. These experiments show the prominent superiority of our CCD-ViT-Small in improving image quality.

**Algorithm 2:** The text pseudo-label  $M_{pl}$  selection.

---

Denote a clustering result  $\Theta \in \{0, 1\}^{H \times W}$ .  
Define  $\Theta_{i,j}$  is the pixel value of row  $i$ , column  $j$ .  
Calculate the sum of pixel values for each side:  
 $L = \sum_{i=1}^H \Theta_{i,1}$ ;  $R = \sum_{i=1}^H \Theta_{i,W}$ ;  
 $T = \sum_{j=1}^W \Theta_{1,j}$ ;  $B = \sum_{j=1}^W \Theta_{H,j}$ ;  
Get condition  $\Gamma$ :  
 $\Gamma = \mathbb{1}_{[T \geq \frac{W}{2}]} + \mathbb{1}_{[B \geq \frac{W}{2}]} + \mathbb{1}_{[L \geq \frac{H}{2}]} + \mathbb{1}_{[R \geq \frac{H}{2}]}$ ,  
**if**  $\Gamma \geq 3$ : **then**  
|  $M_{pl} = 1 - \Theta$   
**else**  
|  $M_{pl} = \Theta$   
**end**

---

**Text segmentation.** In the second row of Table 6, our method yields state-of-the-art results and surpasses other methods by a large margin on downstream text segmentation task. Specifically, CCD-ViT-Small surpasses the Scratch-ViT-Small by 6.7% in terms of Intersection over Union (IoU). Compared with DiG, the first self-supervised learning method to evaluate text segmentation performance on text images, our method shows its prominent superiority, with a performance improvement of 1.7% IoU. These experimental results demonstrate our self-supervised method can learn more generalized text feature representations.

## 5. Ablations and analysis

**Selection of text pseudo-labels.** The clustering result provided by the  $K$ -means algorithm may correspond to a text region or a background region. Thus, as illustrated in Algo.2, we make a minor adaptation to select appropriate text regions from the clustering results by leveraging the observation that text regions are typically located in the center of most scene text instance images, with the four sides of the image being mainly background regions. Compared with randomly selected clustering results ( $\Theta$  or  $1 - \Theta$ ), this adaptation results  $M_{pl}$  leads to a 29.9% IoU performance gain (70.0% vs. 40.1%), as shown in Table 7.

**Effectiveness of self-supervised text segmentation.** In

Table 8. Feature representation evaluation of CCD on scene text recognition benchmarks.

Method	IIIT	SVT	IC13	IC15	SVTP	CUTE	COCO	CTW	TT	HOST	WOST	Avg.
Gen-ViT-Small	86.6	82.1	88.7	72.9	74.4	72.2	48.5	64.1	63.3	33.8	56.5	59.3
Dis-ViT-Small	92.6	90.4	93.4	81.2	81.7	84.0	60.0	72.8	73.1	33.3	56.1	67.0
DiG-ViT-Small	94.2	93.0	95.3	84.3	86.1	87.5	63.4	77.9	75.8	41.7	64.0	<b>71.1</b>
CCD-ViT-Small	93.5	89.6	92.8	82.7	85.1	83.0	60.4	73.3	73.4	47.6	66.5	69.9

Table 9. Comparison results when training with different data ratios.

Label Fraction	Method	IIIT	SVT	IC13	IC15	SVTP	CUTE	COCO	CTW	TT	HOST	WOST	Avg.
1%(27.8K)	DiG-ViT-Small	88.4	86.2	89.9	79.0	76.6	77.8	54.8	67.9	67.2	33.2	53.3	62.9
	CCD-ViT-Small	89.3	86.5	88.8	76.5	80.1	74.7	54.9	65.5	67.8	38.4	55.9	<b>63.7</b>
10%(278K)	DiG-ViT-Small	95.3	94.4	95.9	85.3	87.9	91.7	67.1	80.5	81.1	42.1	64.0	73.5
	CCD-ViT-Small	95.9	94.1	96.6	87.1	89.9	94.1	69.2	81.6	84.3	63.4	76.2	<b>78.2</b>
100%(2.78M)	DiG-ViT-Small	97.7	96.1	97.3	88.6	91.6	96.2	75.0	86.3	88.9	56.0	75.7	80.7
	CCD-ViT-Small	98.0	96.4	98.3	90.3	92.7	98.3	76.7	86.5	91.3	77.3	86.0	<b>84.9</b>

the first row of Table 7, our self-supervised text segmentation network achieves a 3.6% IoU improvement (73.6% vs. 70.0%) by utilizing the text regions clustered by  $K$ -means as pseudo-labels. Additionally, when  $K$ -means is employed for character-to-character representation learning, the average accuracy understandably decreases by 0.24% in row B of Table 10. This can be attributed to two reasons: 1) Neural network has the ability to learn generality from massive training data, which alleviates the noise in the pseudo-labels introduced by  $K$ -means. 2) The underlying morphological representations of glyphs that we need are relatively invariant to slight structural changes, *e.g.*, thicker or thinner, which reduces the dependence on pixel-level high-precision segmentation with expensive costs.

**Feature representation evaluation.** Referring to DiG, we freeze the encoder and train the decoder with ARD. As shown in Table 8, our result is 1.2% lower than DiG (71.1% vs 69.9%). However, this result is not closely related to the final STR results. As evidenced in Table 2 and 3 of DiG, although the discriminative task “Dis- $*$ ” outperforms the generative task “Gen- $*$ ” by 7.7%, it only shows comparable performance to “Gen- $*$ ” in the final STR results. Besides, the reason behind DiG’s superior result is attributed to their utilization of both “Gen- $*$ ” and “Dis- $*$ ”, which brings 4.1% gains. Therefore, when we fairly compare to the same type of pretext task “Dis- $*$ ”, CCD gets a gain of 2.9% (69.9% vs 67.0%). Even compared to DiG, Ours improve by 5.9% (47.6% vs 41.7%) and 2.5% (66.5% vs 64.0%) on the challenging occluded datasets HOST and WOST, respectively.

**Fine-tuning with different data ratios.** To demonstrate the effectiveness of our proposed CCD, we further fine-tuned our method with 1%, 10% and 100% of ARD. As shown in Table 9, CCD outperforms DiG by 0.8%, 4.7% and 4.2%, respectively.

**Effectiveness of distillation strategy.** The distillation strategy serves as a fundamental component in self-supervised

representation learning. To demonstrate the effectiveness of the distillation strategy, we consider CCD without distillation as the baseline model. Specifically, we add the result as shown in row A of Table 10. Compared to Baseline, the total improvement is 4.12%.

**Effectiveness of augmentation strategy.** In Table 10, we compare two augmentation strategies used to establish character-to-character representation consistency across regular and irregular views. 1) “R2R”: both the regular and irregular views adopt color-based augmentations; 2) “CCD”: the default setting, as regular views adopt color-based augmentations, and irregular views adopt a combination of both color- and geometry-based augmentations.

Compared with “R2R” method, the latter achieves 0.49% improvements, which demonstrates large geometric augmentations are still suitable for text images with a diversity of word length, and encourage CCD to reach a data-efficient self-supervised learning regime.

**More comparison results.** 1) We first compared with the language-aware PARSeq (0.5% gains) in Table 4 when training with STD. Besides, we also conduct the comparative experiment using real data provided by PARSeq, resulting in 0.35% gains (+0, +.4, +.3, +.2, +.5, +.6, +.7, +1.0, respectively). 2) Compared to using the cropped patches, CCD achieves 0.7%, 1.1%, 0.0%, 0.7%, 4.3% and 2.0% gains on six standard benchmarks, respectively.

**Further discussion on data-hungry nature of models.** Generally, existing methods achieve significant improvements when training on ARD compared to STD. However, the amount of ARD (only 2.78M) pales in comparison to massive and readily available unlabeled data. Therefore, CCD provides an effective solution by leveraging the intrinsic qualities of unlabeled data (15.77M) to extract robust feature representations, which can generalize well on multiple text-related downstream tasks.

**Further discussion on density-based spatial clustering.**



Table 10. Ablation results on scene text recognition benchmarks.

Index	Method	IIIT	SVT	IC13	IC15	SVTP	CT	Avg.
A	Baseline	95.0	92.9	94.9	85.2	86.7	88.9	91.45
B	<i>K</i> -means	97.6	96.2	98.0	90.5	92.3	97.4	95.33
C	R2R	97.8	96.4	97.3	89.5	92.7	96.4	95.08
E	CCD	98.0	96.4	98.3	90.3	92.7	98.3	<b>95.57</b>

Table 11. Quantification results for different types of clusters.

Types	background	string	semi-character	character
Number	1583	991	1802	<b>14665</b>

The density-based spatial clustering method may encounter difficulties in separating interconnected characters into isolated clusters and has a proneness to cluster an individual character into multiple clusters due to the different densities of connected regions as shown in Fig. 4. Such failures are typically observed in images with tightly interconnected strokes and can also arise due to inaccuracies in text mask predictions generated by our self-supervised text segmentation network. To provide a detailed statistical illustration, we count the number of different types of clusters evaluated on the TextSeg dataset in Table 11. The result was obtained through 7 hours of meticulous manual counting. Clearly, these failures are rare while most clusters still belong to complete characters.

For further discussion, self-supervised learning objective is to obtain general text feature representations, *e.g.*, text foreground features (b) or even character instance features (c) in Fig. 5, while the latter takes a step forward than the former, towards semantic features (d) extracted by supervised learning regimes. As presented above, CCD can obtain more complete characters, enjoying approximately 77% pairwise character alignment under flexible augmentations, towards learning character instance feature representations (c). Even in situations where it struggles to cluster characters due to the issues mentioned above, representation consistency in string or semi-character (both belonging to foreground compared with an item of sequence) can still effectively facilitate the learning of text foreground features (b). Overall, our method can at least degenerate into the same representation consistency learning with sequence-level self-supervised methods in worse-case scenarios, which demonstrates the feasibility and effectiveness of our CCD. More visualization examples about  $M_{pt}$ ,  $M_{seg}$ , and  $S_{reg}$  are shown in Supplementary Material.

**Hyper-parameters.** The density-based spatial clustering method is sensitive to the parameters, namely, *eps* and *min\_samples*. The *eps* controls the granularity of the clustering, while the *min\_samples* sets a threshold for the minimum number of points required for a cluster to be formed. To achieve optimal clustering performance, we conduct a grid search over a range of parameter values and select the




Figure 4. Representative clustering visualization results.



Figure 5. Three different types of text feature representations.




Figure 6. Ablation study of the hyper-parameters.

best combination based on the IoU metric, as illustrated in Fig. 6. Specifically, we first apply this clustering method to the text masks in the training set of Textseg to obtain character clustering results, and then calculate the IoU between these results and the annotated character structures.

## 6. Conclusion

In this paper, we propose a novel self-supervised text recognition method in character-level, termed CCD, which ensures character-to-character representation consistency under flexible augmentations by keeping their pairwise alignment of character regions. Different from existing sequence-to-sequence self-supervised learning models, CCD takes the delineated character structures as basic items for representation learning, and proposes an effective augmentation strategy to enrich the diversity of local character regions. Eventually, CCD shows significant improvement in the robustness and generalizability of the extracted feature representations and refreshes state-of-the-art performance on three text-related tasks.

**Acknowledgements** This work was supported by 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] Aviad Aberdam, Ron Litman, Shahar Tsiper, Oron Anshel, Ron Slossberg, Shai Mazor, R. Manmatha, and Pietro Perona. Sequence-to-sequence contrastive learning for text recognition. In *CVPR*, pages 15302–15312, June 2021. 1, 2, 6
- [2] Yuki M Asano, Christian Rupprecht, and Andrea Vedaldi. A critical analysis of self-supervision, or what we can learn from a single image. *arXiv preprint arXiv:1904.13132*, 2019. 2, 4
- [3] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwal-suk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *ICCV*, pages 4715–4723, 2019. 6
- [4] Darwin Bautista and Rowel Atienza. Scene text recognition with permuted autoregressive sequence models. In *ECCV*, volume 13688, pages 178–196, 2022. 6
- [5] Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, and Yi-Zhe Song. Towards the unseen: Iterative text recognition by distilling from errors. In *ICCV*, pages 14950–14959, October 2021. 2
- [6] Ayan Kumar Bhunia, Shuvojit Ghose, Amandeep Kumar, Pinaki Nath Chowdhury, Aneeshan Sain, and Yi-Zhe Song. Metaht: Towards writer-adaptive handwritten text recognition. In *CVPR*, pages 15830–15839, 2021. 2
- [7] 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
- [8] 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
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9630–9640, 2021. 2, 3, 4, 5
- [10] Jingye Chen, Bin Li, and Xiangyang Xue. Scene text telescope: Text-focused scene image super-resolution. In *CVPR*, pages 12026–12035, 2021. 7
- [11] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, volume 119, pages 1597–1607, 2020. 1, 2, 4
- [12] Da Cheng, Peng Wang, and Cong Yao. Levenshtein ocr. In *ECCV*, volume 13688, pages 322–338, 2022. 6
- [13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE TPAMI*, 38(2):295–307, 2016. 7
- [14] 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. 3
- [15] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, pages 226–231, 1996. 4
- [16] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. In *CVPR*, pages 7098–7107, 2021. 1, 2, 6
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Ávila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap your own latent - A new approach to self-supervised learning. In *NeurIPS*, 2020. 5
- [18] Tongkun Guan, Chaochen Gu, Changsheng Lu, Jingzheng Tu, Qi Feng, Kaijie Wu, and Xinping Guan. Industrial scene text detection with refined feature-attentive network. *IEEE TCSVT*, 32(9):6073–6085, 2022. 1
- [19] Tongkun Guan, Chaochen Gu, Jingzheng Tu, Xue Yang, Qi Feng, Yudi Zhao, and Wei Shen. Self-supervised implicit glyph attention for text recognition. In *CVPR*, pages 15285–15294, 2023. 1, 6
- [20] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *CVPR*, pages 2315–2324, 2016. 1, 5
- [21] 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. 4
- [22] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 15979–15988. IEEE, 2022. 2
- [23] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020. 2
- [24] 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, 6
- [25] Olivier Henaff. Data-efficient image recognition with contrastive predictive coding. In *ICML*, pages 4182–4192, 2020. 2, 4
- [26] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*, 2014. 1, 5
- [27] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *ICDAR*, pages 1156–1160. IEEE, 2015. 5
- [28] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles

- Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. In *ICDAR*, pages 1484–1493. IEEE, 2013. [5](#)
- [29] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, pages 105–114, 2017. [7](#)
- [30] 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 *CVPR Workshops*, pages 546–547, 2020. [1](#), [2](#), [5](#)
- [31] 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](#)
- [32] 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](#), [2](#)
- [33] Ron Litman, Oron Anshel, Shahar Tsiper, Roei Litman, Shai Mazor, and R Manmatha. Scatter: selective context attentional scene text recognizer. In *CVPR*, pages 11962–11972, 2020. [2](#)
- [34] Hao Liu, Bin Wang, Zhimin Bao, Mobai Xue, Sheng Kang, Deqiang Jiang, Yinsong Liu, and Bo Ren. Perceiving stroke-semantic context: Hierarchical contrastive learning for robust scene text recognition. In *AAAI*, pages 1702–1710, 2022. [1](#), [2](#), [3](#), [6](#)
- [35] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*. OpenReview.net, 2019. [5](#)
- [36] Canjie Luo, Lianwen Jin, and Jingdong Chen. Siman: exploring self-supervised representation learning of scene text via similarity-aware normalization. In *CVPR*, pages 1039–1048, 2022. [6](#)
- [37] 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. [2](#)
- [38] Moushumi Medhi, Shubham Sinha, and Rajiv Ranjan Sahay. A text recognition augmented deep learning approach for logo identification. In *Computer Vision, Graphics, and Image Processing - ICVGIP*, volume 10481, pages 145–156, 2016. [1](#)
- [39] Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. In *BMVC*, pages 1–11, 2012. [5](#)
- [40] Ben Niu, Weilei Wen, Wenqi Ren, Xiangde Zhang, Lianping Yang, Shuzhen Wang, Kaihao Zhang, Xiaochun Cao, and Haifeng Shen. Single image super-resolution via a holistic attention network. In *ECCV*, volume 12357, pages 191–207, 2020. [7](#)
- [41] Shuai Peng, Liangcai Gao, Ke Yuan, and Zhi Tang. Image to latex with graph neural network for mathematical formula recognition. In *ICDAR*, volume 12822, pages 648–663, 2021. [1](#)
- [42] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *ICCV*, pages 569–576, 2013. [5](#)
- [43] 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. [2](#), [6](#)
- [44] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications*, 41(18):8027–8048, 2014. [5](#)
- [45] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE TPAMI*, 39(11):2298–2304, 2016. [2](#)
- [46] Amanpreet Singh, Vivek Natarjan, Meet Shah, Yu Jiang, Xinlei Chen, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *CVPR*, pages 8317–8326, 2019. [1](#)
- [47] Amanpreet Singh, Guan Pang, Mandy Toh, Jing Huang, Wojciech Galuba, and Tal Hassner. Textocr: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text. In *CVPR*, pages 8802–8812, 2021. [5](#)
- [48] 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](#)
- [49] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *ICCV*, pages 1457–1464. IEEE, 2011. [5](#)
- [50] Peng Wang, Cheng Da, and Cong Yao. Multi-granularity prediction for scene text recognition. In *ECCV*, volume 13688, pages 339–355, 2022. [1](#), [6](#)
- [51] 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. [1](#)
- [52] Wenjia Wang, Enze Xie, Xuebo Liu, Wenhao Wang, Ding Liang, Chunhua Shen, and Xiang Bai. Scene text image super-resolution in the wild. In *ECCV*, volume 12355, pages 650–666, 2020. [5](#), [7](#)
- [53] 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. [2](#), [6](#)
- [54] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004. [5](#)
- [55] Jiajia Wu, Jun Du, Fengren Wang, Chen Yang, Xinzhe Jiang, Jinshui Hu, Bing Yin, Jianshu Zhang, and Lirong Dai. A multimodal attention fusion network with a dynamic vocabulary for textvqa. *PR*, 122:108214, 2022. [1](#)
- [56] Xudong Xie, Ling Fu, Zhifei Zhang, Zhaowen Wang, and Xiang Bai. Toward understanding wordart: Corner-guided transformer for scene text recognition, 2022. [2](#)

- [57] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: a simple framework for masked image modeling. In *CVPR*, pages 9643–9653, 2022. [2](#)
- [58] Xingqian Xu, Zhifei Zhang, Zhaowen Wang, Brian Price, Zhonghao Wang, and Humphrey Shi. Rethinking text segmentation: A novel dataset and a text-specific refinement approach. In *CVPR*, pages 12045–12055, 2021. [5](#)
- [59] Ruijie Yan, Liangrui Peng, Shanyu Xiao, and Gang Yao. Primitive representation learning for scene text recognition. In *CVPR*, pages 284–293, 2021. [6](#)
- [60] 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. [2](#)
- [61] Mingkun Yang, Minghui Liao, Pu Lu, Jing Wang, Sheng-gao Zhu, Hualin Luo, Qi Tian, and Xiang Bai. Reading and writing: Discriminative and generative modeling for self-supervised text recognition. In *ACM MM*, pages 4214–4223, 2022. [1](#), [2](#), [5](#), [6](#), [7](#)
- [62] Zhengyuan Yang, Yijuan Lu, Jianfeng Wang, Xi Yin, Dinei Florêncio, Lijuan Wang, Cha Zhang, Lei Zhang, and Jiebo Luo. TAP: text-aware pre-training for text-vqa and text-caption. In *CVPR*, pages 8751–8761, 2021. [1](#)
- [63] 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. [2](#), [6](#)
- [64] Xiaoyu Yue, Zhanghui Kuang, Chenhao Lin, Hongbin Sun, and Wayne Zhang. Robustscanner: Dynamically enhancing positional clues for robust text recognition. In *ECCV*, pages 135–151. Springer, 2020. [2](#)
- [65] Jinglei Zhang, Tiancheng Lin, Yi Xu, Kai Chen, and Rui Zhang. Relational contrastive learning for scene text recognition, 2023. [2](#)
- [66] 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. [1](#), [6](#)
- [67] Cairong Zhao, Shuyang Feng, Brian Nlong Zhao, Zhijun Ding, Jun Wu, Fumin Shen, and Heng Tao Shen. Scene text image super-resolution via parallelly contextual attention network. In *ACM MM*, pages 2908–2917, 2021. [7](#)
- [68] Dajian Zhong, Shujing Lyu, Palaiahnakote Shivakumara, Bing Yin, Jiajia Wu, Umapada Pal, and Yue Lu. Sgbnet: Semantic gan and balanced attention network for arbitrarily oriented scene text recognition. In *ECCV*, 2022. [2](#), [6](#)
- [69] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan L. Yuille, and Tao Kong. Image BERT pre-training with online tokenizer. In *ICLR*, 2022. [2](#)