

# Conditional Text Image Generation with Diffusion Models

Yuanzhi Zhu, Zhaohai Li, Tianwei Wang, Mengchao He, Cong Yao<sup>†</sup>  
Alibaba DAMO Academy, Hangzhou, China.

{z.yuanzhi, zhaohai.li, wangtw}@foxmail.com, mengchao.hmc@alibaba-inc.com,  
yaocong2010@gmail.com.

## Abstract

Current text recognition systems, including those for handwritten scripts and scene text, have relied heavily on image synthesis and augmentation, since it is difficult to realize real-world complexity and diversity through collecting and annotating enough real text images. In this paper, we explore the problem of text image generation, by taking advantage of the powerful abilities of Diffusion Models in generating photo-realistic and diverse image samples with given conditions, and propose a method called **Conditional Text Image Generation with Diffusion Models (CTIG-DM for short)**. To conform to the characteristics of text images, we devise three conditions: image condition, text condition, and style condition, which can be used to control the attributes, contents, and styles of the samples in the image generation process. Specifically, four text image generation modes, namely: (1) synthesis mode, (2) augmentation mode, (3) recovery mode, and (4) imitation mode, can be derived by combining and configuring these three conditions. Extensive experiments on both handwritten and scene text demonstrate that the proposed CTIG-DM is able to produce image samples that simulate real-world complexity and diversity, and thus can boost the performance of existing text recognizers. Besides, CTIG-DM shows its appealing potential in domain adaptation and generating images containing *Out-Of-Vocabulary (OOV)* words.

## 1. Introduction

Text recognition has been an important research topic in the computer vision community for a long time, due to its wide range of applications. In the past few years, numerous recognition methods for scene and handwritten text [3, 16, 43, 57, 58, 65, 69, 72] have been proposed, which have substantially improved the recognition accuracy on various benchmarks. The volume and diversity of data are crucial for high recognition performance, but it is extremely

<sup>†</sup>Corresponding author.




Figure 1. Handwritten text image samples from IAM [46] or generated by our proposed CTIG-DM. On the left, the handwriting styles of the same word “and” written by different writers are considerably different, indicating the diversity of handwritten text and the challenge of handwritten text recognition. On the right, two images out of four in each row are written by the corresponding writer on the left. Can you distinguish them from the generated samples? (The answer will be revealed in the next page.)



Figure 2. Scene text image samples from Real-L [5] or produced by CTIG-DM. Only seven images herein are real. Can you identify them? (The answer will be revealed in the next page.)

hard, if not impossible, to collect and label sufficient real text images, so majority of the existing recognition methods rely heavily on data synthesis and augmentation.

Previously, a variety of data synthesis and augmentation methods [7, 17, 18, 20, 25, 33, 38, 44, 45, 67] have been proposed to enrich data for training stronger text recognition models. In this paper, we investigate a technique that is highly related and complementary to such works. Drawing inspiration from the recent progress of Diffusion Models [15, 48], we propose a text image generation model, which is able to conduct data synthesis, and thus can boost

the performance of existing text recognizers.

A recent study [15] has shown that State-Of-The-Art (SOTA) likelihood-based models [48] can outperform GAN-based methods [8, 30, 68] in generating images. Diffusion models [23, 48, 59] have been becoming increasingly popular, due to their powerful generative ability in various vision tasks [2, 10, 11, 41, 55]. A typical representative of diffusion models is Denoising Diffusion Probabilistic Models (DDPM) [23]. It generates diverse samples through different initial states of simple distribution and each transition. This means that it is challenging for DDPM to control the content of the output image due to the randomness of the initial states and transitions. Guided-Diffusion [15] provides conditions to diffusion models by adding classifier guidance. UnCLIP [53] further pre-trains a CLIP model [52] to match the image and whole text, which are used as the conditions for the diffusion models in image generation. While these approaches have focused on natural images, images with handwritten or scene text have their unique characteristics (as shown in Fig. 1 and Fig. 2), which require not only *image fidelity and diversity*, but also *content validity* of the generated samples, i.e., the text contained in the images should be the *same* as specified in the given conditions.

In this paper, we present a diffusion model based conditional text image generator, termed **Conditional Text Image Generation with Diffusion Models (CTIG-DM)** for short. To the best of our knowledge, this is one of the first works to introduce diffusion models into the area of text image generation. The proposed CTIG-DM consists of a conditional encoder and a conditional diffusion model. Specifically, the conditional encoder generates three conditions, i.e., image condition, text condition, and style condition (the writing style of a specific writer). These conditions are proved to be critical for the fidelity and diversity of the generated text images. The conditional diffusion part uses these conditions to generate images from random Gaussian noise. As can be seen in Fig. 1 and Fig. 2, the quality of the images generated by CTIG-DM is quite high that one can hardly tell them from real images\*. By combining the given conditions, four image generation modes can be derived, i.e., synthesis mode, augmentation mode, recovery mode, and imitation mode. With these modes, various text images that can be used to effectively boost the accuracy of existing text recognizers (see Sec. 4 for more details) could be produced. Moreover, CTIG-DM shows its potential in handling OOV image generation and domain adaptation.

The contributions can be summarized as follows:

- We propose a text image generation method based on diffusion models, which is one of the first attempts to use diffusion models to generate text images.

\*In Fig. 1, the real images are the first and last of each row. In Fig. 2, the real images are even numbered.

- We devise three conditions and four image generation modes, which can facilitate the generation of text images with high validity, fidelity, and diversity.
- Experiments on both scene text and handwritten text demonstrate that CTIG-DM can significantly improve both the image quality and the performance of previous text recognizers. Besides, CTIG-DM is effective in OOV image generation and domain adaptation.

## 2. Related Work

### 2.1. Text Recognition

As an important task in computer vision, text recognition has attracted extensive attention in the community. Specifically, Scene Text Recognition (STR) and Handwritten Text Recognition (HTR) are the most popular research directions [9, 37, 75].

Scene text images generally contain complex backgrounds and irregular text arrangements. Early, He et al. [21] and Shi et al. [57] proposed to model STR as a sequence-to-sequence mapping issue by combining CNN, RNN, and CTC [19]. Then, attention-based methods [43, 58, 65] gradually emerged and achieved a breakthrough in irregular text recognition. In recent years, benefiting from the success of Transformer [61], many methods [12, 16, 64, 69] improved the recognizer from the perspective of the language model.

Handwritten text images have diverse writing styles and difficult-to-segment cursive joins. Zhang et al. [74] addressed the handwriting style diversity problem by domain adaption. Bhunia et al. [3] employ Model Agnostic meta-learning algorithm to train writer adaptive HTR network. Recently, due to the lack of real data, more researchers paid attention to the fields of text data augmentation and synthesis, thereby improving the performance of handwritten text recognizers [1, 6, 17, 27, 44, 45].

### 2.2. Text Image Augmentation and Synthesis

Wigington et al. [67] and Bhunia et al. [6] built grids on the original images and hidden features, respectively, and then augmented them by adding random perturbations. Luo et al. [45] proposed a learnable augmentation method to obtain more controllable samples. These methods have made significant progress in improving the performance of text recognizers. However, they fail to create OOV samples and thus the diversity is limited by the training set.

There are many GAN-based handwritten text image synthesis methods. Fogel et al. [17] presented a semi-supervised approach that can generate images of words with a variable length. Kang et al. [27] generated credible handwritten word images by adjusting calligraphic style features and textual contents. Luo et al. [44] proposed a style bank

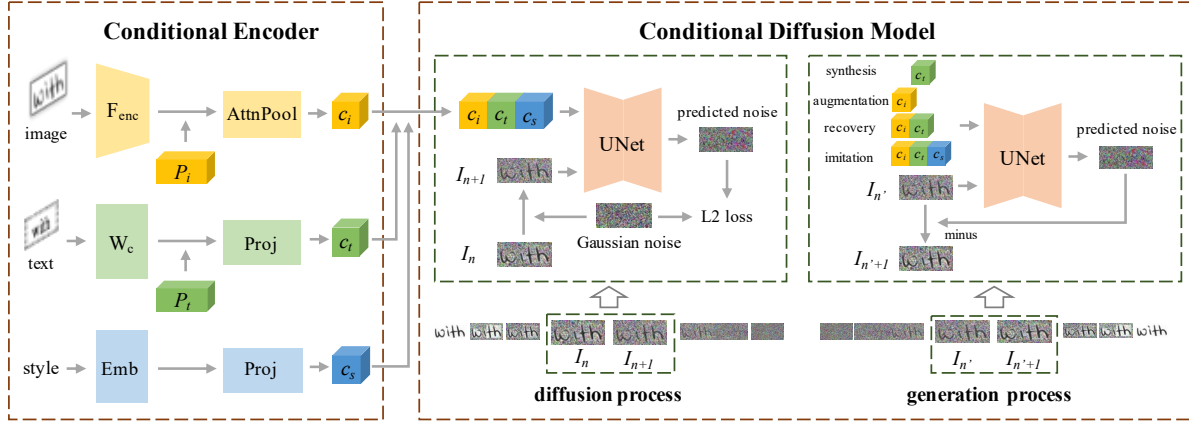


Figure 3. The overall architecture of the proposed CTIG-DM.

and dual discriminators to solve the issues of style representation and content embedding. However, GAN-based methods are generally difficult to train and require careful design of hyper-parameters, otherwise, it is easy to fall into mode collapse [15]. Besides, most of the above data synthesis methods are only experimented on handwritten Latin text without exploring generalization to other types of texts, such as scene text or handwritten Chinese text.

### 2.3. Diffusion Models

Recently, a category of deep generative models, named diffusion models [59], has achieved impressive results in computer vision tasks, outperforming GAN-based methods in the diversity of generated images [15]. Inspired by the non-equilibrium thermodynamics theory [59], Ho et al. [23] learned to model the Markov transition from noise to data distribution, enabling unconditional image generation. Luhman et al. [42] focused on online handwriting and generated point sequences based on diffusion models. Prafulla et al. [15] introduced additional classifiers to provide conditions for the diffusion models. Nichol et al. [49] explored diffusion models for text-conditional image synthesis with classifier-free guidance. Ramesh et al. [53] proposed a two-stage model, i.e., a prior that generates a CLIP image embedding given a text caption, and a decoder that generates an image conditioned on the image embedding.

Unlike general image generation, text image generation requires more unique contexts and textural features at the character-level. This inspires us to use a text recognizer to obtain conditions related to text images.

## 3. Methodology

### 3.1. Conditional Encoder

As illustrated in Fig. 3, the outputs of the conditional encoder consist of an image condition  $c_i$ , a text condition  $c_t$ , and a style condition  $c_s$ . Specifically,  $c_i$  represents the unique visual characteristics of input images.  $c_t$  indicates the semantic characteristics of input contexts.  $c_s$  describes

the information about the writer styles, which are only used in the handwritten text generation.

**Image condition.** Different from the patterns of natural images, the rich visual information of text images concentrates on the characters. Therefore, we propose to use the encoder of a pre-trained text recognizer to obtain image condition that can better express *general features* (e.g., textural and colors) rather than the noisy features (e.g. the backgrounds). The background information comes from the training data and is encoded in the trained diffusion models. Given the input image  $I$ , the generation process of the image condition can be formulated as

$$c_i = AttnPool(F_{enc}(I) + Emb(P_i)), \quad (1)$$

where  $F_{enc}$  is the feature extractor of the pre-trained text recognizer.  $P_i$  denotes the index of encoded image patches, and  $Emb$  is the embedding function which encodes  $P_i$  to obtain position embedding.  $AttnPool$  represents the attention pooling [34] to customize visual representations. The image condition is an image-level representation, which is an aggregation of patch-level features.

**Text condition.** Text condition specifies the *contents* of the generated text images and represents the unique contexts among chars, which is critical for the proposed CTIG-DM. Although the pre-trained word embedding [14] is widely used in the natural language processing field, it cannot handle the text image generation of OOV words. Therefore, we adapt the classifier weights  $W_c$  of the pre-trained text recognizer described above to encode text condition. The generation process of text condition can be described as

$$c_t = Proj(W_c T + Emb(P_t)), \quad (2)$$

where  $T$  and  $P_t$  represent the one-hot encoding and the index of characters in the text string label. A linear projection layer  $Proj$  is applied to unify the dimensionality with the text condition.

**Style condition.** The style condition is particularly designed for text generation of HTR, which contains the writing characteristics (e.g., character slants, cursive joins, and

stroke widths) of a specific writer. In other words, the style condition represents *personal style* rather than image style. The generation of style condition can be represented as

$$c_s = Proj(Emb(S)), \quad (3)$$

where  $S$  is the writer ID. Overall, the combination of  $c_i$ ,  $c_t$ , and  $c_s$  is fed into the conditional diffusion model to generate text images.

### 3.2. Conditional Diffusion Model

Different from natural image generation in vanilla diffusion models [15, 23, 48], the proposed CTIG-DM introduces expert knowledge related to text image, i.e., image condition, text condition, and style condition. As shown in Fig. 3, following [23], the conditional diffusion model is implemented by UNet [56], which contains two processes, i.e., the diffusion process and the generation process.

**Diffusion process.** In the diffusion process, by continuously adding Gaussian noise to an initial image, the characteristics of the initial image will gradually disappear, and the image eventually becomes standard Gaussian noise. Specifically, as illustrated in Fig. 3, at the diffusion step  $n$ , the next noisy image  $I_{n+1}$  can be obtained by

$$I_{n+1} = I_n + \epsilon, \quad (4)$$

where  $\epsilon$  denotes the Gaussian noise. Then the next noisy image and the proposed conditions are fed into the UNet to predict the distribution of the noise, which is supervised by the squared error loss

$$L = \|\epsilon - \epsilon_\theta([c_i, c_t, c_s], I_{n+1})\|^2. \quad (5)$$

$\epsilon_\theta$  and  $[\cdot]$  represent the parameters of UNet and the concatenation process respectively. Specifically, given the diffusion step,  $L$  can be directly calculated by the original image rather than the intermediate status [23].

**Generation process.** During the generation process, the image is finally generated by gradually denoising the random initial Gaussian noise. Specifically, at the generation step  $n'$ , the current denoising image  $I_{n'}$  and proposed conditions are fed into UNet to predict the noise. Then the next denoising image  $I_{n'+1}$  is the minus between predicted noise and the current denoising image, as follows,

$$I_{n'+1} = I_{n'} - \epsilon_\theta(\mathbf{c}, I_{n'}). \quad (6)$$

$\mathbf{c}$  stands for the different settings of proposed conditions. Benefiting from the input of various conditions, we propose four image generation modes, i.e., synthesis mode, augmentation mode, recovery mode, and imitation mode. Tab. 1 shows the details for each generation mode. In empirical experiments, the text condition determines the diversity and is helpful for OOV image generation, while the image condition and style condition influence the fidelity. More details can be found in Sec. 4.3.

Table 1. The combination of conditions for each generation mode.

Generation Modes	$c_i$	$c_t$	$c_s$
Synthesis		✓	
Augmentation	✓		
Recovery	✓	✓	
Imitation	✓	✓	✓

## 4. Experiments

### 4.1. Datasets

**Handwritten text datasets.** IAM [46] contains more than 115,000 words written in English by 657 different writers. RIMES [32] contains more than 60,000 words written in French by over 1,000 authors. IAM and RIMES are widely used in previous methods [6, 17, 44, 45, 60] and can serve for a variety of handwritten recognition tasks. CVL [32] contains seven different handwritten texts (one in German and six in English) written by 311 different writers. We use the English part for the experiment of domain adaptation. CASIA-HWDB 1.0-1.1 [36] consists of 2,678,424 images of offline handwritten Chinese characters. We use it for OOV handwritten Chinese characters generation.

**Scene text datasets.** We use MJSynth [24, 25], SynthText [20], and Real-L [5], as training data. The test datasets consist of regular datasets, i.e., IIIT 5K-Words (IIIT) [47], Street View Text (SVT) [63], ICDAR 2003 (IC03) [40] and ICDAR 2013 (IC13) [29], and irregular datasets i.e., ICDAR 2015 (IC15) [28], Street View Text-Perspective (SVTP) [50] and CUTE80 (CUTE) [54]. Details of these datasets can be found in previous works [69].

### 4.2. Implementation Details

**Handwritten text.** Similar to previous work [17, 44] on the IAM, RIMES, and CVL datasets, the height of the training images is resized to 64 pix and the width is calculated with the original aspect ratio (up to 256 pix). The evaluation criteria of recognition performance are Word Error Rate (WER %) and Character Error Rate (CER %). The WER respects the ratio of the error at the word level, and the CER corresponds to the edit distance between the recognition result and ground-truth, normalized by the length of ground-truth. With respect to the generation quality, FID [22], GS [31], SSIM, RMSE and LPIPS [73] are introduced to our experiments. Lower values of WER, CER, FID, GS, RMSE and LPIPS, and higher value of SSIM are preferable.

**Scene text.** When combined with previous methods, all experimental settings are kept the same for the sake of fair comparison, except that generated data from the proposed CTIG-DM is used.

**Network.** The pre-train text recognizer uses CRNN architecture [57] and the diffusion models follow the DDPM architecture [23, 48]. Specifically, the proposed conditions, whose dimensionality is set to 512, are concatenated to the

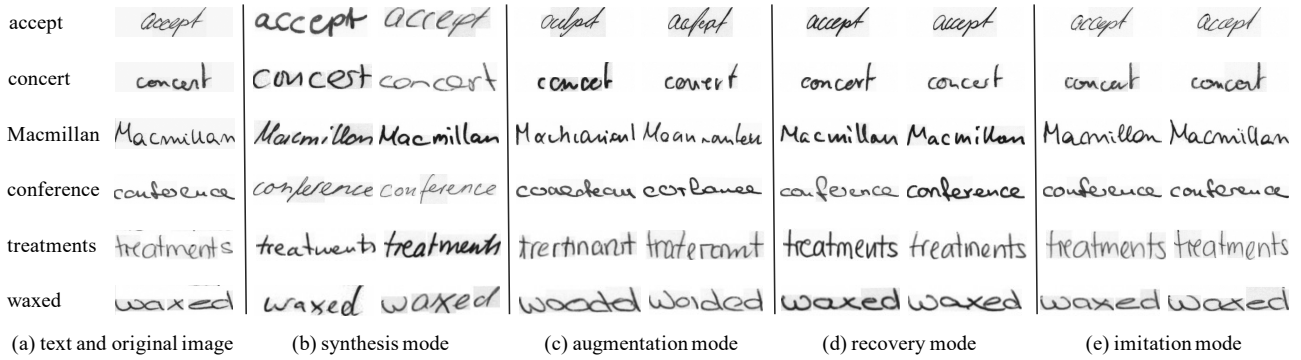


Figure 4. Visualization of handwritten text images and their text strings and different generation modes.

time-step embedding described in GLIDE [49]. To avoid the proposed CTIG-DM over-relying on any condition, we randomly set the  $c_i$ ,  $c_t$ , and  $c_s$  to a learnable embedding 20%, 10%, and 20%, respectively, of the time.

**Optimization.** We use the AdamW [39] as the optimizer with the settings of  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ , and  $weight\_decay = 0.2$ . The learning rate is 0.0001 with cosine annealing. The batch size is set to 64 for training recognizer and 256 for training diffusion models. All experiments are conducted on NVIDIA Tesla V100 GPUs.

### 4.3. Ablation Study

In this subsection, we conduct ablation study on the IAM dataset. Concretely, we explore the role of various conditions obtained by the conditional encoder in the diffusion process and show visualizations of each generation mode during the generation process.

As presented in Tab. 2, the trends of different metrics are consistent. First, the baseline is built without any conditions, which performs a large FID score of 33.42. Then, we add  $c_i$  and  $c_t$  respectively, and find that both FID scores are significantly improved, which indicates that the image and text features have critical supports for text image generation. Besides, by comparing the above two experiments, we find that  $c_t$  has a greater impact on the quality of image generation. When both  $c_i$  and  $c_t$  are added, the FID further decreases to 9.76. Moreover, we add  $c_s$  to further guide the proposed CTIG-DM to present specific handwriting styles that can be easily identified. Finally, with all conditions, our method achieves the best FID of 9.34.

Table 2. Effectiveness of different conditions in diffusion process.

$c_i$	$c_t$	$c_s$	FID ↓	SSIM ↑	RMSE ↓	LPIPS ↓
			33.42	0.5346	0.2195	0.4144
✓			13.36	0.6476	0.1994	0.2337
	✓		10.92	0.7077	0.1738	0.1349
✓	✓		9.76	0.7241	0.1697	0.1132
✓	✓	✓	9.34	0.7374	0.1501	0.0936

Benefiting from the input of diverse conditions, Fig. 4 and Fig. 5 illustrate the image visualizations of different generation modes for handwritten text and scene text, re-

spectively. Specifically, since the synthesis mode only depends on the  $c_t$ , we observe that its generated images have rich diversity. As presented in Fig. 4 (b) and Fig. 5 (b), handwritten text images exist in a variety of character slants, ink blots, cursive joins, stroke widths, and paper backgrounds, while scene text images are diverse in text rotations, backgrounds, blur noise, and fonts. Correspondingly, the  $c_i$  is more critical for the fidelity of the generated images. In augmentation mode, we find that the generated images are similar in overall appearance to the original images, but part details are lost in the specific characters. As shown in Fig. 4 (c) and Fig. 5 (c), images are more likely to generate wrong characters when the  $c_i$  is used solely. This is because only the image attributes are included in  $c_i$  and the unique contexts among characters included in  $c_t$  are missing. After adding  $c_t$ , the generated images in Fig. 4 (d) and Fig. 5 (d) rarely have wrong characters and contain diversity at the same time. Finally, when  $c_s$  designed for handwritten text is added, the fidelity of the generated images in Fig. 4 (e) improves further, indicating that  $c_s$  plays an important role in image generation. In Sec. 4.6, we will demonstrate the effect of  $c_s$  on the style control of generated text images. Since the synthesis mode and the imitation mode (recovery mode in scene text) have the best diversity and fidelity, respectively, the generated data (denoted as MIX) we use in Sec. 4.4 are derived from an equal-proportion mixture of these two modes.

### 4.4. Recognition Performance

For text recognition tasks, the ultimate purpose of generating images is to augment the training set and improve the performance of the recognizers. Therefore, we conduct experiments on multiple types of text recognition, including English handwritten text recognition, French handwritten text recognition, and scene text recognition, to demonstrate the validity of the generated data on the recognizers.

#### 4.4.1 Handwritten Text Recognition

Following the settings of previous works [6, 44, 45], we use CRNN [57] as the recognizer. Fig. 6 illustrates the effec-



Figure 5. Visualization of scene text images and their text strings and different generation modes.

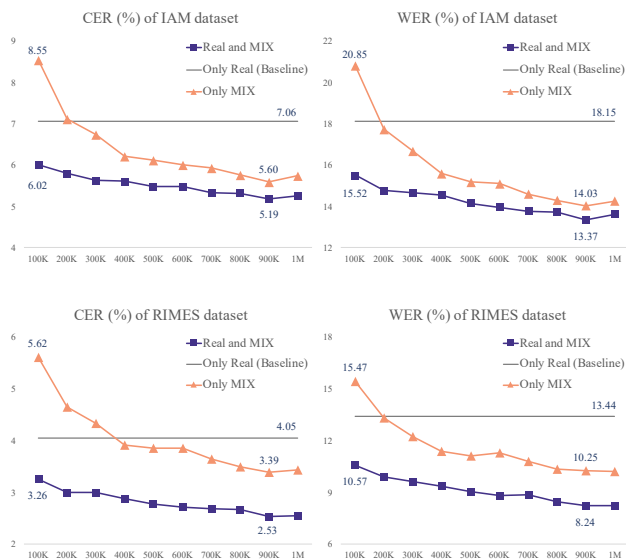


Figure 6. Effectiveness of our generated data (denoted as MIX) on IAM and RIMES (denoted as Real). The horizontal and vertical axes represent the amount of MIX used and the recognition performance (lower values are preferable), respectively.

tiveness of our generated data. The horizontal line on each subfigure indicates the baseline where the real training set is used solely. Specifically, for IAM dataset, about 200K generated data can achieve comparable performance as the baseline. Simultaneously, as the amount of MIX gradually increases, CER and WER are steadily decreasing, and realize the best results, i.e., CER of 5.60% and WER of 14.03%, at 900K. The experiment results on RIMES dataset show similar trends. Moreover, as described in the purple line in Fig. 6, the recognizer performance can be further improved when using both real and generated data. Finally, compared with the baseline, on the IAM dataset, WER and CER decrease by 4.78% and 1.87%, respectively, and on the RIMES dataset, WER and CER decrease by 5.20% and 1.52%, respectively. This indicates that in HTR with a limited real training set, the generated data of our method can significantly improve the recognizer performance, which demonstrates the validity of the proposed CTIG-DM.

Table 3. Comparison of recognition performance with previous methods on the IAM and RIMES datasets. The numbers of “ $\Delta$ ” denote the improvements from the baseline to our method. “Aug\*” represents the random geometric augmentation of [45].

	IAM		RIMES	
	WER	CER	WER	CER
Sueiras et al. [60]	23.80	8.80	15.90	4.80
Alonso et al. [1]	-	-	11.90	4.03
Zhang et al. [74]	22.20	8.50	-	-
Bhunia et al. [6]	17.19	8.41	10.47	6.44
ScrabbleGAN [17]	23.61	13.42	11.32	3.57
Kang et al. [27]	17.26	6.75	-	-
Luo et al. [45]	14.04	5.34	9.23	2.57
SLOGAN [44]	14.97	5.95	11.50	3.35
Baseline	18.15	7.06	13.44	4.05
Ours	<b>13.37</b>	<b>5.19</b>	<b>8.24</b>	<b>2.53</b>
$\Delta$	<b>+4.78</b>	<b>+1.87</b>	<b>+5.20</b>	<b>+1.52</b>
Ours + Aug*	<u>12.01</u>	<u>4.67</u>	<u>6.89</u>	<u>1.98</u>

We compare our method to SOTA methods in Tab. 3. For a fair comparison, note that methods using additional real data or language models are outside the scope of this study. It can be seen that the proposed CTIG-DM outperforms previous methods for data augmentation and data synthesis. Moreover, we integrate our method with data augmentation by using the open-source toolkit<sup>†</sup> and performing random geometric augmentation [45] on the generated samples. Finally, adding data augmentation to our method can further improve the performance, which suggests that the proposed CTIG-DM is complementary to previous works.

#### 4.4.2 Scene Text Recognition

Different STR methods may adopt different backbones, data processing, training policies, etc. Therefore, for a fair comparison, we reproduce two representative works CRNN and ABINet to evaluate the effectiveness of our algorithm. As shown in Tab. 4 (a), when combined with our generated data, the performance of CRNN is improved by an average of 1.7% on the six test datasets and that of ABINet is

<sup>†</sup><https://github.com/Canjie-Luo/Text-Image-Augmentation>

Table 4. Comparison of recognition performance with previous methods on STR datasets. MIX represents our generated data and we note its amount in parentheses. The numbers of “ $\Delta$ ” in **green** and **blue** denote the improvements over each dataset and average, respectively.

	Methods	Training Dataset	IC13	SVT	IIIT	IC15	SVTP	CUTE	Average
	CRNN [57]	ST	86.7	80.8	78.2	-	-	-	-
	TBRA [4]	MJ+ST	93.6	87.5	87.9	77.6	79.2	74.0	84.6
	ESIR [71]	MJ+ST	91.3	90.2	93.3	76.9	79.6	83.3	87.1
	MORAN [43]	MJ+ST	92.4	88.3	91.2	68.8	76.1	77.4	83.3
	ASTER [58]	MJ+ST	91.8	89.5	93.4	76.1	78.5	79.5	86.4
	SAM [35]	MJ+ST	95.3	90.6	93.9	77.3	82.2	87.8	88.3
	SE-ASTER [51]	MJ+ST	92.8	89.6	93.8	80.0	81.4	83.6	88.3
	TextScanner [62]	MJ+ST	92.9	90.1	93.9	79.4	84.3	83.3	88.5
	DAN [65]	MJ+ST	93.9	89.2	94.3	74.5	80.0	84.4	87.2
	RobustScanner [70]	MJ+ST	94.8	88.1	95.3	77.1	79.5	90.3	88.4
	SRN [69]	MJ+ST	95.5	91.5	94.8	82.7	85.1	87.8	90.4
	VisionLAN [66]	MJ+ST	95.7	91.7	95.8	83.7	86.0	88.5	91.2
	ABINet [16]	MJ+ST	97.4	93.5	96.2	86.0	89.3	89.2	92.6
(a)	CRNN (Baseline)	ST	87.8	80.9	85.7	63.0	66.3	70.8	77.2
	CRNN (Ours)	ST+MIX (15M)	89.7	81.6	87.1	65.0	67.8	75.0	78.9
	$\Delta$	-	<b>+1.9</b>	<b>+0.7</b>	<b>+1.4</b>	<b>+2.0</b>	<b>+1.5</b>	<b>+4.2</b>	<b>+1.7</b>
	ABINet (Baseline)	MJ+ST	97.3	93.5	96.3	86.1	89.1	89.2	92.7
	ABINet (Ours)	MJ+ST+MIX (30M)	<b>98.1</b>	<b>94.2</b>	<b>96.6</b>	<b>86.7</b>	<b>89.6</b>	<b>91.7</b>	<b>93.2</b>
	$\Delta$	-	<b>+0.8</b>	<b>+0.7</b>	<b>+0.3</b>	<b>+0.6</b>	<b>+0.5</b>	<b>+2.5</b>	<b>+0.5</b>
(b)	CRNN (Baseline)	Real-L	86.6	77.1	84.0	61.1	62.7	64.2	75.0
	CRNN (Ours)	Real-L+MIX (2M)	89.0	81.1	88.1	67.8	68.8	72.2	79.9
	$\Delta$	-	<b>+2.4</b>	<b>+4.0</b>	<b>+4.1</b>	<b>+6.7</b>	<b>+6.1</b>	<b>+8.0</b>	<b>+4.9</b>
	ABINet (Baseline)	Real-L	95.7	93.5	96.6	86.9	86.5	94.1	92.8
	ABINet (Ours)	Real-L+MIX (2M)	97.2	96.0	97.0	89.0	90.2	95.8	94.3
	$\Delta$	-	<b>+1.5</b>	<b>+2.5</b>	<b>+0.4</b>	<b>+2.1</b>	<b>+3.7</b>	<b>+1.7</b>	<b>+1.5</b>

boosted by an average of 0.5%, which is considerable.

By comparing Tab. 3 and Tab. 4 (a), we observe that the improvements of CTIG-DM in STR are not as obvious as that of HTR. This is probably because most STR methods are trained on synthesized data only, and the training set has about 14M images, which is two orders of magnitude larger. Both the diversity and fidelity of the proposed CTIG-DM are partially limited by the synthesized training set. However, CTIG-DM can still bring performance improvements to the advanced recognizer without making any changes to the network. Further, we evaluate the effectiveness of our method on the real dataset Real-L, whose scale is about 270K. As shown in Tab. 4 (b), our generated data brings significant improvements of 4.9% and 1.5% to CRNN and ABINet, respectively. It proves that the proposed CTIG-DM is able to produce valid image samples that simulate the complexity and diversity of the real world, which can better improve the performance of existing text recognizers.

#### 4.5. Quality of Generated Images

In this subsection, we compare with previous methods in terms of image generation quality. Similar to the settings in [44], FID is calculated with 25k real and 25k generated images, while GS is calculated with 5k real and 5k gener-

Table 5. Comparison of image generation quality with previous methods on HTR datasets. Lower values are preferable.

	RIMES		IAM	
	FID	GS	FID	GS
Alonso et al. [1]	23.94	$8.58 \times 10^{-4}$	-	-
ScrabbleGAN [17]	23.78	$7.60 \times 10^{-4}$	20.72	$2.56 \times 10^{-2}$
HiGAN [18]	-	-	17.28	-
CG-GAN [33]	-	-	19.03	-
HWT [7]	-	-	19.40	$1.01 \times 10^{-2}$
Davis et al. [13]	23.72	$7.19 \times 10^{-1}$	20.65	$4.88 \times 10^{-2}$
SLOGAN [44]	12.06	$5.59 \times 10^{-4}$	-	-
Ours	<b>9.75</b>	<b><math>6.23 \times 10^{-5}</math></b>	<b>9.34</b>	<b><math>5.91 \times 10^{-5}</math></b>

ated images. As illustrated in Tab. 5, benefiting from the proposed various condition and powerful generative ability of the diffusion models, our method outperforms the SOTA methods by a notable margin in all metrics on both datasets, which indicates the advancement of CTIG-DM.

#### 4.6. Generating Out-of-Vocabulary Images

As described in Sec. 3.2, the proposed synthesis mode only depends on the text condition. Therefore, by changing the input text, we can generate images containing OOV text. In this subsection, we explore the OOV image generation quality of the proposed CTIG-DM. Since previous works [7, 26, 33, 44] generate OOV images with the spec-

ified styles, for a fair comparison, we use text condition and style condition as the input conditions. As presented in Tab. 6, compared with previous methods, the proposed CTIG-DM significantly boosts the OOV image generation quality. Fig. 7 shows the visualization of images generated by CTIG-DM, which suggests the controllability of our method over the styles and text of the generated images.

Table 6. Comparison of OOV image generation quality with previous methods on the IAM dataset.

Method	FID
GANwriting [26]	125.87
HWT [7]	109.45
SLOGAN [44]	97.81
CG-GAN [33]	104.81
Ours	<b>25.52</b>

Styles	Generated images with OOV text
<i>touch</i>	<i>Similar cloud heroic nodding</i>
<i>walked</i>	<i>deated Married lick warrior</i>
<i>adult</i>	<i>thunder Push card planets</i>
<i>great</i>	<i>Clever senders receive mood</i>

Figure 7. Visualization of images generated by CTIG-DM according to the specified styles and OOV text. Note that none of the words on the right appear in the training set.

#### 4.7. Domain Adaptation

Table 7. Comparison of domain adaptation capacity with previous methods on the CVL dataset. SYN represents our generated data using the lexicon of the CVL training set.

Method	Training Data	WER	CER
Baseline	IAM	42.64	18.49
ScrabbleGAN [17]	IAM+GAN (100K)	35.98	17.27
SLOGAN [44]	IAM+GAN (100K)	34.98	14.10
Ours	IAM+SYN (100K)	<b>26.24</b>	<b>10.89</b>

In this subsection, we investigate the domain adaptation capacity of our method. Following the settings of the previous works [17, 44], we first train CTIG-DM with the IAM training set and then generate 100K samples using the lexicon of the CVL training set. Finally, combining the IAM training set and generated samples to train the recognizer, we evaluate the performance on the test set of CVL. We repeat the training five times and report the averages. As illustrated in Tab. 7, the performance improvements brought by our generated data far exceed that of the previous methods. Concretely, compared with the baseline, WER and CER are decreased by 16.40% and 7.60%, respectively, indicating the superior validity and diversity of our generated data.



Figure 8. Visualization of (a) non-existing handwritten Chinese characters (radical sequences are presented in the first column) and (b) ancient rare characters.

#### 4.8. Applications

The above experiments already show that CTIG-DM can advance Latin text image generation. In this subsection, we explore the applications of CTIG-DM in other scenarios, e.g., the generation of OOV handwritten Chinese and ancient characters, to illustrate the potential of our method.

**OOV handwritten Chinese characters generation.** Compared with Latin characters such as English and French, Chinese characters are more complex in structure and contain more information. Therefore, in order to make full use of the Chinese character prior, we decompose the Chinese characters into radical sequences as text conditions. Similar to generating OOV words in English, we can construct radical sequences that are not in the training set to generate OOV Chinese characters or create non-existing handwritten Chinese characters. Fig. 8 (a) shows some visualizations where the characters are not in the training set or are even nonexistent. Specifically, the characters shown in the last line are typos, which may come into play in scenarios where typo recognition is required.

**Ancient characters generation.** Ancient text recognition has great significance in the inheritance of historical knowledge, literature, and art. One of its main difficulties is the long-tailed distribution, i.e., it lacks training samples for rare characters. Therefore, we can use CTIG-DM to generate rare characters (shown in Fig. 8 (b)). In addition, since rare characters generally have an extremely small number of samples, the method of combining CTIG-DM and radicals for characters generation is also promising.

#### 5. Conclusion and Future Work

We have presented a text image generation method, called CTIG-DM, based on diffusion models. To adapt to the nature of text images, we devise three conditions and four generation modes. Extensive experiments have demonstrated the effectiveness and advantage of CTIG-DM. Specifically, it can be used to boost the performance of existing text recognizers. In the future, we will explore more generation modes on more types of text images, and generate text images with the styles of unseen writers.



## References

- [1] Eloi Alonso, Bastien Moysset, and Ronaldo O. Messina. Adversarial generation of handwritten text images conditioned on sequences. In *ICDAR*, pages 481–486, 2019. 2, 6, 7
- [2] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *CVPR*, pages 18187–18197, 2022. 2
- [3] Amandeep Kumar Pinaki Nath Chowdhury Aneeshan Sain Yi-Zhe Song Ayan Kumar Bhunia, Shuvojit Ghose. Metaht: Towards writer-adaptive handwritten text recognition. In *CVPR*, 2021. 1, 2
- [4] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwalsuk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *ICCV*, pages 4714–4722, 2019. 7
- [5] Jeonghun Baek, Yusuke Matsui, and Kiyoharu Aizawa. What if we only use real datasets for scene text recognition? toward scene text recognition with fewer labels. In *CVPR*, pages 3113–3122, 2021. 1, 4
- [6] Ayan Kumar Bhunia, Abhirup Das, Ankan Kumar Bhunia, Perla Sai Raj Kishore, and Partha Pratim Roy. Handwriting recognition in low-resource scripts using adversarial learning. In *CVPR*, pages 4762–4771, 2019. 2, 4, 5, 6
- [7] Ankan Kumar Bhunia, Salman Khan, Hisham Cholakkal, Rao Muhammad Anwer, Fahad Shahbaz Khan, and Mubarak Shah. Handwriting transformers. In *ICCV*, pages 1086–1094, 2021. 1, 7, 8
- [8] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 2
- [9] Xiaoxue Chen, Lianwen Jin, Yuanzhi Zhu, Canjie Luo, and Tianwei Wang. Text recognition in the wild: A survey. *ACM Comput. Surv.*, 54(2):42:1–42:35, 2022. 2
- [10] Jooyoung Choi, Sungwon Kim, Yonghyun Jeong, Youngjune Gwon, and Sungroh Yoon. ILVR: conditioning method for denoising diffusion probabilistic models. In *ICCV*, pages 14347–14356, 2021. 2
- [11] Hyungjin Chung, Byeongsu Sim, and Jong Chul Ye. Come-closer-diffuse-faster: Accelerating conditional diffusion models for inverse problems through stochastic contraction. In *CVPR*, pages 12403–12412, 2022. 2
- [12] Cheng Da, Peng Wang, and Cong Yao. Levenshtein OCR. In *ECCV*, pages 322–338, 2022. 2
- [13] Brian L. Davis, Bryan S. Morse, Brian L. Price, Chris Tensmeyer, Curtis Wigington, and Rajiv Jain. Text and style conditioned GAN for the generation of offline-handwriting lines. In *BMVC*, 2020. 7
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186, 2019. 3
- [15] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, pages 8780–8794, 2021. 1, 2, 3, 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, 7
- [17] Sharon Fogel, Hadar Averbuch-Elor, Sarel Cohen, Shai Mazon, and Roei Litman. ScrabbleGAN: Semi-supervised varying length handwritten text generation. In *CVPR*, pages 4323–4332, 2020. 1, 2, 4, 6, 7, 8
- [18] Ji Gan and Weiqiang Wang. Higan: Handwriting imitation conditioned on arbitrary-length texts and disentangled styles. In *AAAI*, pages 7484–7492, 2021. 1, 7
- [19] Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML*, pages 369–376, 2006. 2
- [20] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *CVPR*, pages 2315–2324, 2016. 1, 4
- [21] Pan He, Weilin Huang, Yu Qiao, Chen Change Loy, and Xiaoou Tang. Reading scene text in deep convolutional sequences. In *AAAI*, pages 3501–3508, 2016. 2
- [22] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NIPS*, pages 6626–6637, 2017. 4
- [23] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2, 3, 4
- [24] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. In *NIPS Deep Learning Workshop*, 2014. 4
- [25] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Reading text in the wild with convolutional neural networks. *IJCV*, 116(1):1–20, 2016. 1, 4
- [26] Lei Kang, Pau Riba, Yaxing Wang, Marçal Rusiñol, Alicia Fornés, and Mauricio Villegas. GANwriting: Content-conditioned generation of styled handwritten word images. In *ECCV*, pages 273–289, 2020. 7, 8
- [27] Lei Kang, Marçal Rusiñol, Alicia Fornés, Pau Riba, and Mauricio Villegas. Unsupervised writer adaptation for synthetic-to-real handwritten word recognition. In *WACV*, pages 3502–3511, 2020. 2, 6
- [28] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman K. Ghosh, Andrew D. Bagdanov, and Masakazu Iwamura. ICDAR 2015 competition on robust reading. In *ICDAR*, pages 1156–1160, 2015. 4
- [29] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, and Sergi Robles Mestre. ICDAR 2013 robust reading competition. In *ICDAR*, pages 1484–1493, 2013. 4
- [30] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020. 2

- [31] Valentin Khrulkov and Ivan V. Oseledets. Geometry score: A method for comparing generative adversarial networks. In *ICML*, pages 2626–2634, 2018. 4
- [32] Florian Kleber, Stefan Fiel, Markus Diem, and Robert Sablatnig. CVL-database: An offline database for writer retrieval, writer identification and word spotting. In *ICDAR*, pages 560–564, 2013. 4
- [33] Yuxin Kong, Canjie Luo, Weihong Ma, Qiyuan Zhu, Sheng-gao Zhu, Nicholas Yuan, and Lianwen Jin. Look closer to supervise better: One-shot font generation via component-based discriminator. In *CVPR*, pages 13472–13481, 2022. 1, 7, 8
- [34] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam R. Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*, pages 3744–3753, 2019. 3
- [35] Minghui Liao, Pengyuan Lyu, Minghang He, Cong Yao, Wenhao Wu, and Xiang Bai. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. *IEEE TPAMI*, 43(2):532–548, 2021. 7
- [36] Cheng-Lin Liu, Fei Yin, Da-Han Wang, and Qiu-Feng Wang. CASIA online and offline chinese handwriting databases. In *ICDAR*, pages 37–41, 2011. 4
- [37] Shangbang Long, Xin He, and Cong Yao. Scene text detection and recognition: The deep learning era. *International Journal of Computer Vision*, 129:161–184, 2018. 2
- [38] Shangbang Long and Cong Yao. Unrealtext: Synthesizing realistic scene text images from the unreal world. In *CVPR*, pages 5488–5497, 2020. 1
- [39] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 5
- [40] Simon M. Lucas, Alex Panaretos, Luis Sosa, Anthony Tang, Shirley Wong, and Robert Young. ICDAR 2003 robust reading competitions: entries, results, and future directions. *Int. J. Document Anal. Recognit.*, 7(2-3):105–122, 2005. 4
- [41] Andreas Lugmayr, Martin Danelljan, Andrés Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, pages 11451–11461, 2022. 2
- [42] Troy Luhman and Eric Luhman. Diffusion models for handwriting generation. *CoRR*, abs/2011.06704, 2020. 3
- [43] Canjie Luo, Lianwen Jin, and Zenghui Sun. MORAN: A multi-object rectified attention network for scene text recognition. *PR*, 90:109–118, 2019. 1, 2, 7
- [44] Canjie Luo, Yuanzhi Zhu, Lianwen Jin, Zhe Li, and Dezhi Peng. SLOGAN: Handwriting style synthesis for arbitrary-length and out-of-vocabulary text. *IEEE TNNLS*, pages 1–13, 2022. 1, 2, 4, 5, 6, 7, 8
- [45] Canjie Luo, Yuanzhi Zhu, Lianwen Jin, and Yongpan Wang. Learn to augment: Joint data augmentation and network optimization for text recognition. In *CVPR*, pages 13743–13752, 2020. 1, 2, 4, 5, 6
- [46] U-V Marti and Horst Bunke. The IAM-database: an english sentence database for offline handwriting recognition. *Int. J. Document Anal. Recognit.*, 5(1):39–46, 2002. 1, 4
- [47] Anand Mishra, Karteek Alahari, and C. V. Jawahar. Scene text recognition using higher order language priors. In *BMVC*, pages 1–11, 2012. 4
- [48] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, pages 8162–8171, 2021. 1, 2, 4
- [49] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, and Bob McGrew. GLIDE: towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, pages 16784–16804, 2022. 3, 5
- [50] 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. 4
- [51] Zhi Qiao, Yu Zhou, Dongbao Yang, Yucan Zhou, and Weiping Wang. SEED: Semantics enhanced encoder-decoder framework for scene text recognition. In *CVPR*, pages 13528–13537, 2020. 7
- [52] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, and Sandhini Agarwal. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 2
- [53] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents, 2022. 2, 3
- [54] Anhar Risnumawan, Palaiahnakote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Syst. Appl.*, 41(18):8027–8048, 2014. 4
- [55] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10674–10685, 2022. 2
- [56] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 4
- [57] 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. 1, 2, 4, 5, 7
- [58] Baoguang Shi, Mingkun Yang, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. ASTER: An attentional scene text recognizer with flexible rectification. *IEEE TPAMI*, 2019. 1, 2, 7
- [59] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, pages 2256–2265, 2015. 2, 3
- [60] Jorge Sueiras, Victoria Ruíz, Ángel Sánchez, and José Francisco Vélez. Offline continuous handwriting recognition using sequence to sequence neural networks. *Neurocomputing*, 289:119–128, 2018. 4, 6
- [61] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017. 2
- [62] Zhaoyi Wan, Minghang He, Haoran Chen, Xiang Bai, and Cong Yao. TextScanner: Reading characters in order for robust scene text recognition. In *AAAI*, number 07, pages 12120–12127, 2020. 7

- [63] Kai Wang, Boris Babenko, and Serge J. Belongie. End-to-end scene text recognition. In *ICCV*, pages 1457–1464, 2011. [4](#)
- [64] Peng Wang, Cheng Da, and Cong Yao. Multi-granularity prediction for scene text recognition. In *ECCV*, pages 339–355, 2022. [2](#)
- [65] 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*, 2020. [1](#), [2](#), [7](#)
- [66] 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 14174–14183, 2021. [7](#)
- [67] Curtis Wigington, Seth Stewart, Brian L. Davis, Bill Barrett, Brian L. Price, and Scott Cohen. Data augmentation for recognition of handwritten words and lines using a CNN-LSTM network. In *ICDAR*, pages 639–645, 2017. [1](#), [2](#)
- [68] Yan Wu, Jeff Donahue, David Balduzzi, Karen Simonyan, and Timothy Lillicrap. Logan: Latent optimisation for generative adversarial networks. *arXiv preprint arXiv:1912.00953*, 2019. [2](#)
- [69] Deli Yu, Xuan Li, Chengquan Zhang, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text recognition with semantic reasoning networks. In *CVPR*, pages 12113–12122, 2020. [1](#), [2](#), [4](#), [7](#)
- [70] 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, 2020. [7](#)
- [71] Fangneng Zhan and Shijian Lu. ESIR: End-to-end scene text recognition via iterative image rectification. In *CVPR*, pages 2059–2068, 2019. [7](#)
- [72] Jianshu Zhang, Jun Du, and Lirong Dai. Multi-scale attention with dense encoder for handwritten mathematical expression recognition. In *International Conference on Pattern Recognition*, pages 2245–2250, 2018. [1](#)
- [73] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, pages 586–595, 2018. [4](#)
- [74] Yaping Zhang, Shuai Nie, Wenju Liu, Xing Xu, Dongxiang Zhang, and Heng Tao Shen. Sequence-to-sequence domain adaptation network for robust text image recognition. In *CVPR*, pages 2740–2749, 2019. [2](#), [6](#)
- [75] Yingying Zhu, Cong Yao, and Xiang Bai. Scene text detection and recognition: recent advances and future trends. *Frontiers of Computer Science*, 10:19–36, 2016. [2](#)