

Scene Text Image Super-resolution based on Text-conditional Diffusion Models

Chihiro Noguchi Shun Fukuda Masao Yamanaka
Toyota Motor Corporation, Japan

{chihiro_noguchi_aa, shun_fukuda, masao_yamanaka}@mail.toyota.co.jp

Abstract

Scene Text Image Super-resolution (STISR) has recently achieved great success as a preprocessing method for scene text recognition. STISR aims to transform blurred and noisy low-resolution (LR) text images in real-world settings into clear high-resolution (HR) text images suitable for scene text recognition. In this study, we leverage text-conditional diffusion models (DMs), known for their impressive text-to-image synthesis capabilities, for STISR tasks. Our experimental results revealed that text-conditional DMs notably surpass existing STISR methods. Especially when texts from LR text images are given as input, the text-conditional DMs are able to produce superior quality super-resolution text images. Utilizing this capability, we propose a novel framework for synthesizing LR-HR paired text image datasets. This framework consists of three specialized text-conditional DMs, each dedicated to text image synthesis, super-resolution, and image degradation. These three modules are vital for synthesizing distinct LR and HR paired images, which are more suitable for training STISR methods. Our experiments confirmed that these synthesized image pairs significantly enhance the performance of STISR methods in the TextZoom evaluation.

1. Introduction

Scene text recognition has attracted considerable attention because of its high applicability in many areas. However, it remains a challenging task when the input images are significantly blurred and low-resolution. To address these problems, scene text image super-resolution (STISR) is a promising approach. STISR aims to restore the high-resolution (HR) text images from low-resolution (LR) text images. Unlike standard single-image super-resolution [10, 23–26, 60], which is applied to general scene images, STISR is specialized for text images. Hence, the essence of STISR lies in restoring HR text images while preserving the textual information contained in the LR images. The restored images are called super-resolution (SR) images.

To effectively train STISR models, there is a need for a



Figure 1. SR text images restored from different methods. Text below the images indicates text recognition results. Red characters indicate incorrect or missing results. TCDM denotes the text-conditional diffusion model, while TCDM* indicates TCDM trained using ground-truth text input.

substantial quantity of paired LR-HR text images. However, obtaining such paired images from real-world scenarios is significantly expensive, which results in a scarcity of public datasets tailored for STISR. An alternative strategy is the generation of synthetic LR images. This can be achieved by applying synthetic degradation, such as bicubic interpolation and blur kernels, to text images. However, the synthetic images produced via these simple degradation methods might not be apt for real-world settings, where more complex degradations can occur [51].

In this study, we aim to produce realistic paired text images that are better suited for training STISR methods. The proposed framework in this study hinges on two primary elements. Firstly, it utilizes text-conditional diffusion models. Diffusion models (DMs) [15, 45, 46] have demonstrated remarkable generative performance in various tasks, especially in text-to-image synthesis [34, 37, 40]. This generative prowess can also be considered effective for STISR. Unlike text-to-image synthesis settings, in STISR, text prompts are not predetermined. Therefore, a pretrained text recogni-

tion model is required to extract text features from the text images. Our experimental results indicate that while the vanilla DM is sufficiently effective, incorporating text feature extraction further elevates its performance. The second element arises from the observation that the performance of STISR methods can be substantially boosted by using ground-truth texts of LR text images, rather than depending on outputs of the text recognition model. These two elements allow our framework to generate high-quality super-resolution text images from the provided text images and their corresponding ground-truth texts.

Furthermore, our proposed method integrates two additional components. The first is an image degradation model that generates LR text images from the provided text images. Combined with the aforementioned super-resolution model, our framework can generate LR-HR paired text images from the input text images with their ground-truth texts. The second component is a text image generation model that generates text images from text strings. Once these models are trained, they allow for the generation of as many LR-HR paired images as desired. In our proposed framework, all the three components are implemented using text-conditional DMs.

Synthesising LR-HR paired images is not novel strategy to enhance single image super-resolution performance. BSRGAN [56] and Real-ESRGAN [53] combine different types of synthetic degradations to emulate realistic ones. Sidiya et al. [4] introduce two Style-GAN-based models [22] for synthesizing paired image. MCinCGAN [59] and Pseudo-SR [30], which are based on Cycle-GAN [63], eliminate the need for image-level paired training images. However, these methods are not specialized for text images. Conversely, our proposed framework is designed specifically for text images, enabling the generation of high-quality paired images by utilizing the textual content in the text images as input.

The super-resolution and image degradation models in our framework are trained using TextZoom [51], the most commonly dataset for STISR. For the text image generation model, we explore two scenarios: one where the model is trained exclusively on TextZoom and another where it is trained on both TextZoom and additional text image datasets. The paired images synthesized by our framework are combined with TextZoom to form augmented datasets. Our experiments reveal that the augmented datasets enhance the performance of STISR methods, showing a marked improvement over relying solely on TextZoom.

Our contributions can be summarized as follows. (1) To the best of our knowledge, this study represents the pioneering effort applying text-conditional DMs to scene text image super-resolution. Experimental validation confirms that the DMs attain state-of-the-art performance when assessed on TextZoom [51]. (2) We introduce a framework to gener-

ate LR-HR paired text images. Using ground-truth texts of input text images for training each component of the framework, it can generate high-quality paired images that are apt for the training datasets of STISR methods. (3) Through our experiments, we show that the augmented datasets effectively enhance the performance of STISR methods, surpassing results obtained by relying solely on TextZoom.

2. Related Works

2.1. Scene Text Recognition

Scene text recognition methods basically accept a cropped image that contains only a single word as input. The text images are fed into the CNN- or Transformer-based feature extractor. Several decoding methods have been proposed using image features to predict characters. Two major decoding methods exist: CTC decoding [13, 42] and attention decoding [3, 12, 27, 43]. Please refer to the comprehensive survey [13] for more details on scene text recognition.

2.2. Scene Text Image Super-Resolution

The goal of STISR is to restore an HR image from an LR image. Unlike standard single image super-resolution, in STISR, one word always appears in the input image (called a text image). Therefore, extracting accurate textual contents from text images and effectively utilizing it is essential to generate high-quality SR images. In fact, not only PSNR/SSIM metrics, which are commonly used in single image super-resolution, but also text recognition accuracy are used to evaluate the performance of STISR methods.

Many methods have been proposed in this area. Dong et al. [11] applied an SRCNN [10] to obtain SR images and demonstrated its effectiveness in terms of scene text recognition performance. TextSR [52] and MCGAN [54] are based on SRGAN [25] and use loss functions to adopt the guidance of the text recognizer. PCAN [61] and [36] focused on the high-frequency components of an image during reconstruction. PlugNet [32] and TSRN [51] utilized novel modules based on residual blocks to enhance SR quality. In STT [5], position- and content-aware losses were proposed. These losses encourage the model to focus on character regions and labels, which helps produce SR text images suitable for text recognition. TATT [28, 29] incorporates text prior information obtained from a text recognizer into a transformer-based network using cross-attention modules. C3-STISR [62] exploits three clues: text prior, visual, and linguistic information. TG [6] was proposed to concentrate on the stroke-level internal structures of characters. Wang et al. [51] proposed TextZoom, which contains real LR-HR paired images and is widely used to evaluate the performance of STISR methods.

2.3. Diffusion Models

Gaussian DMs were first introduced in [45] and have been improved for image generation [9, 15, 16, 35] and various other tasks [2, 19]. DMs achieve significant performance, particularly in text-to-image synthesis. A major concern here is how to condition the DM to control the image generated using text prompts accurately. To this end, classifier guidance was proposed in [9], which can improve the sample quality using class labels while reducing the diversity of the generated images. Classifier guidance was extended to classifier-free guidance in [17], eliminating the need for an additional classifier. This makes it easier to condition DMs on information that is difficult to handle with a classifier such as text [34, 37, 37, 40]. In addition, DMs have been applied to single image super-resolution [41]. The LR image is created by downsampling the HR image, concatenating it with the input image, and feeding it to the model. Additionally, SR is often used as an auxiliary task in text-to-image synthesis to generate high-resolution images [34].

3. Methodology

The method proposed in this paper is based on denoising diffusion probabilistic models (DDPMs). In the following, we describe the improvement of DDPMs for application to STISR in Sec. 3.1. Loss functions for extracting textual information and the model architecture of the proposed method are described in Sec. 3.1.1 and 3.1.2, respectively. Next, we introduce a framework for synthesizing LR-HR paired text images in Sec. 3.2.

3.1. Text-Conditional DMs for STISR

A DM consists of forward and reverse processes. In the forward process, Gaussian noise is gradually added to the input image, and eventually it becomes pure Gaussian noise. Conversely, in the reverse process, starting from pure Gaussian noise, noise is sequentially removed to recreate the original image.

The objective here is to achieve each step of the reverse process from a Gaussian noise input. Recent successful models view this problem as one of predicting the Gaussian noise contained at each step of the forward process with a sequence of denoising autoencoders $\epsilon_\theta(x_t, t)$, $\forall t \in \{1, \dots, T\}$. Here, x_t is the output of the t th step of the forward process. Consequently, the loss function is given as

$$\mathcal{L}_u = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2], \quad (1)$$

where t is sampled uniformly from $\{1, \dots, T\}$ and $\epsilon \sim \mathcal{N}(0, I)$. Note that x_T corresponds to a noise image. The derivation of Eq. 1 is provided in Sec. A and [15].

From the STISR perspective, it is necessary to control the output HR text image x_0 under two different conditions: LR text image and text in the text image. According to

prior studies on text-to-image synthesis, this can be easily achieved by simply conditioning $\epsilon_\theta(x_t, t)$ in Eq. 1.

When conditioning a DM on an LR image x^l , x_t is replaced with $x'_t = [x_t, x^l]$, where $[\cdot, \cdot]$ indicates the operation of concatenation in the channel dimension. Here, x^l must be bicubic interpolated to the size of x_t .

Unlike text-to-image synthesis, text prompts are not provided by users in STISR. Therefore, textual information must be extracted from the provided text images. In previous studies [28, 29], such textual information, which is referred to as *text prior*, is extracted using a pretrained text recognition method. This model is called *text prior generator*. In this study, we additionally introduce ground-truth text prior, which is created from user-provided text instead of using the text prior generator. The text prior or the ground-truth text prior is then fed into a text encoder to obtain text features z . Consequently, the loss function in Eq. 1 is modified as follows:

$$\mathcal{L}_c = \mathbb{E}_{t, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x'_t, z, t)\|^2]. \quad (2)$$

3.1.1 Text Prior Generator

Following prior studies [28, 29, 62], the text prior generator is trained concurrently with the DM. We use intermediate features from one layer before the final output as text features. As the input to the text prior generator, we can use t th output image x_t in addition to LR image x^l . When t is small, x_t can be sufficiently clear to be useful for the text prior generator, whereas when t is large, x_t may still be noisy. Therefore, we consider the weighted sum of the two images $y_t = (1 - (t/T)^3)x_t + (t/T)^3x^l$ as the input for the text prior generator. When t is large, y_t is close to x^l ; as t decreases, y_t approaches x_t .

When ground-truth text prior is used, a text prior generator is not required. Instead, the ground-truth text prior must be manually created from the corresponding texts and is represented by a matrix $f_P \in \mathbb{R}^{l \times |A|}$, where l is the maximum length of the input text, and A is the character set. Each row of f_P is a one-hot vector. When the text length is less than l , all zero vectors are inserted at equal intervals to fit the size of f_P .

3.1.2 Model Architecture

In our text-conditional DMs, we adopt the prevailing UNet architecture [39] similar to previous studies [9, 35]. However, unlike text-image synthesis, STISR does not require a large language model for the text encoder because the input text is typically only a single word. Therefore, we adopt a simple self-attention-based architecture for the text encoder (see Sec. B for more details).

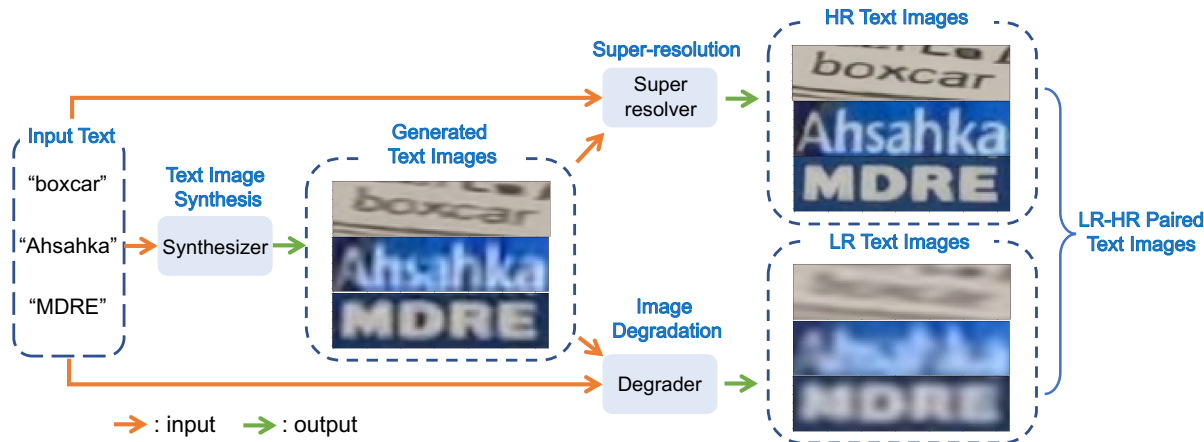


Figure 2. Overview of the proposed framework for LR-HR paired text image synthesis. First, text images are generated by Synthesizer with user-provided text. Next, the generated text images are fed to Super-resolver and Degradation to generate HR and LR images, respectively. Synthesizer, Super-resolver, and Degradation are based on text-conditional DMs, which take user-provided texts as input.

To condition DMs on text features obtained from the text encoder, cross-attention modules [48] have been a common choice in text-to-image synthesis. In previous studies, cross-attention modules are inserted in multiple resolution layers [38, 40]. On the other hand, in our text-conditional DMs, cross-attention modules are inserted only at the bottom of the UNet. Our experimental results show that this architecture can sufficiently grasp the text features while curbing computational expenses. However, a single cross-attention module is inadequate, thus we adopt an architecture with multiple cross-attention modules. See Sec. B for the architecture details.

3.2. LR-HR Paired Text Image Synthesis

In this subsection, we explore the use of text-conditional DMs to build a framework for synthesizing LR-HR paired text images. As shown in Fig. 2, the proposed framework comprises three distinct subtasks: (1) text image synthesis, (2) super-resolution, and (3) image degradation. We refer to the text-conditional DMs for these three tasks as (1) *Synthesizer*, (2) *Super-resolver*, and (3) *Degradation*, respectively. The workflow of the proposed framework is as follows. First, Synthesizer generates text images from user-provided texts. Then, these generated text images are fed to Super-resolver and Degradation to generate HR and LR images, respectively. As a result, from an arbitrary text input, the corresponding LR-HR paired text images can be created. The three text-conditional DMs do not share parameters and depend on different inputs. Synthesizer is solely conditioned on input texts, whereas Super-resolver and Degradation are conditioned both on the text images from Synthesizer and the initial input texts.

Training Synthesizer requires only the text images and

their corresponding texts. In other words, it can be trained using scene text recognition datasets, which are more readily available in large volumes than STISR datasets. It is noteworthy that using Synthesizer is optional. Scene text recognition datasets can be directly used as input to Super-resolver and Degradation, instead of using text images generated by Synthesizer. Regarding the training of Super-resolver and Degradation, both paired LR-HR text images and their corresponding texts are required. The Degradation training can be realized by simply swapping the condition with the target image in the Super-resolver training. For the Degradation training, it is not trivial whether to use texts as input. Nonetheless, our experimental results suggest an enhancement in the quality of the generated text images. Refer to Sec. E for more details.

Incorporating both Super-resolver and Degradation is crucial to our framework. In scene images, text areas frequently occupy small portions. As a result, text image datasets encompass many low-resolution and blurred images as well as high-resolution ones. By applying these two modules to the provided text images, our framework can generate distinct pairs of LR and HR images.

4. Evaluation

4.1. Evaluation on TextZoom

In this subsection, we evaluate the performance of our text-conditional DMs on TextZoom and compare them with state-of-the-art methods in STISR.

Training Details. The size of the HR images was 32×128 , and the LR images were bicubic interpolated to the same size before being fed to the DMs. Following prior studies, we used a character set consisting of digits and low-

Method	CRNN [42]				MORAN [27]				ASTER [43]			
	Easy	Medium	Hard	Avg.	Easy	Medium	Hard	Avg.	Easy	Medium	Hard	Avg.
Bicubic	36.4%	21.1%	21.1%	26.8%	60.6%	37.9%	30.8%	44.1%	67.4%	42.4%	31.2%	48.2%
HR	76.4%	75.1%	64.6%	72.4%	91.2%	85.3%	74.2%	84.1%	94.2%	87.7%	76.2%	86.6%
SRCNN [10]	41.1%	22.3%	22.0%	29.2%	63.9%	40.0%	29.4%	45.6%	70.6%	44.0%	31.5%	50.0%
SRResNest [25]	45.2%	32.6%	25.5%	35.1%	66.0%	47.1%	33.4%	49.9%	69.4%	50.5%	35.7%	53.0%
TSRN [51]	52.5%	38.2%	31.4%	41.4%	70.1%	55.3%	37.9%	55.4%	75.1%	56.3%	40.1%	58.3%
STT [5]	59.6%	47.1%	35.3%	48.1%	74.1%	57.0%	40.8%	58.4%	75.7%	59.9%	41.6%	60.1%
PCAN [61]	59.6%	45.4%	34.8%	47.4%	73.7%	57.6%	41.0%	58.5%	77.5%	60.7%	43.1%	61.5%
TG [6]	61.2%	47.6%	35.5%	48.9%	75.8%	57.8%	41.4%	59.4%	77.9%	60.2%	42.4%	61.3%
TATT [29]	62.6%	53.4%	39.8%	52.6%	72.5%	60.2%	43.1%	59.5%	78.9%	63.4%	45.4%	63.6%
C3-STISR [62]	65.2%	53.6%	39.8%	53.7%	74.2%	61.0%	43.2%	60.5%	79.1%	63.3%	46.8%	64.1%
DDPM*	66.8%	56.5%	41.8%	55.0%	78.4%	62.2%	45.3%	62.0%	81.1%	64.3%	48.9%	64.7%
TCDM	67.3%	57.3%	42.7%	55.7%	77.6%	62.9%	45.9%	62.2%	81.3%	65.1%	50.1%	65.5%

Table 1. Comparison with the existing methods in terms of the recognition accuracy on TextZoom. TCDM denotes the text-conditional DM. * indicates DDPM [35] trained on TextZoom by ourselves.

Method	GT Text	SSIM ($\times 10^{-2}$)	PSNR	Acc. (%)
Bicubic		69.61	20.35	26.8
HR		-	-	72.4
SRCNN [10]		72.27	20.78	29.2
SRResNest [25]		74.03	21.03	35.1
TSRN [51]		76.90	21.42	41.4
STT [5]		76.14	21.05	48.1
PCAN [61]		77.52	21.49	47.4
TG [6]		74.56	21.40	48.9
TATT [29]		79.30	21.52	52.6
C3-STISR [62]		77.21	21.51	53.7
DDPM*		79.50	22.70	55.0
TCDM		79.58	22.83	55.7
TATT	✓	79.34	22.35	61.0
TCDM	✓	80.25	22.86	68.1

Table 2. Comparison with the existing methods in terms of average of SSIM/PSNR and recognition accuracy on TextZoom. The recognition accuracy was evaluated by CRNN. The bottom two lines show the results of the methods using ground-truth texts.

erases letters plus a blank character when training the text prior generator ($|\mathcal{A}| = 37$). When training with ground-truth text prior, it was created with a character set consisting of digits and lower- and upper-case letters plus a black character ($|\mathcal{A}| = 63$). Additionally, we set the maximum length of the input text $l = 26$. To fine-tune the text prior generator, we used the KL loss [28]. Other detailed hyperparameters are provided in Sec. B. All experiments were performed using a workstation equipped with A100 GPUs.

Dataset. To train and evaluate the STISR methods, we used TextZoom [51], which contains 21,740 LR-HR pair images and their corresponding text labels. Images in TextZoom were captured by cameras with different focal lengths in the wild. The training set of TextZoom consisted of 17,367 pairs, and the rest were used as the test set. The test set was divided into three subsets according to the camera focal length: easy (1,619 pairs), medium (1,411 pairs), and hard (1,343 pairs).

Evaluation Metrics. We used three metrics to evaluate the performance of STISR methods. The first two metrics are SSIM and PSNR, which are widely used in single image super-resolution to measure similarity with HR images. The third one is text recognition accuracy. Following prior studies, we used three text recognition methods, CRNN [42], MORAN [27] and ASTER [43] to evaluate the recognition accuracy.

4.1.1 Comparison with State-of-the-art Methods

Table 1 shows the comparison results with existing methods in terms of the recognition accuracy on TextZoom. DDPM corresponds to the vanilla DMs without text conditions. The two DM-based models outperformed the existing methods, and our text-conditional DM achieved an even better performance than DDPM. In addition, Tab. 2 shows the results in terms of the average SSIM/PSNR and recognition accuracy for the three difficulty levels of the TextZoom test sets. Notably, the text-conditional DM outperforms the existing methods. The bottom two rows present the results of TATT and the text-conditional DM when the ground-truth text prior is used. We can see that their performances are significantly improved by using ground-truth texts.

4.2. Evaluation on Augmented Datasets

We present the experimental results to demonstrate the effectiveness of the LR-HR paired text images synthesized by the proposed framework. We augmented TextZoom with the synthesized images and trained our text-conditional DM and the state-of-the-art model TATT [29] on the augmented dataset. The TextZoom test set was used for the performance evaluation. Examples of the generated text images are shown in Fig. 7.

Training Details. For text inputs to our framework, we randomly chose the length of a word from 2 to 13 and then randomly selected an English word of that length from a

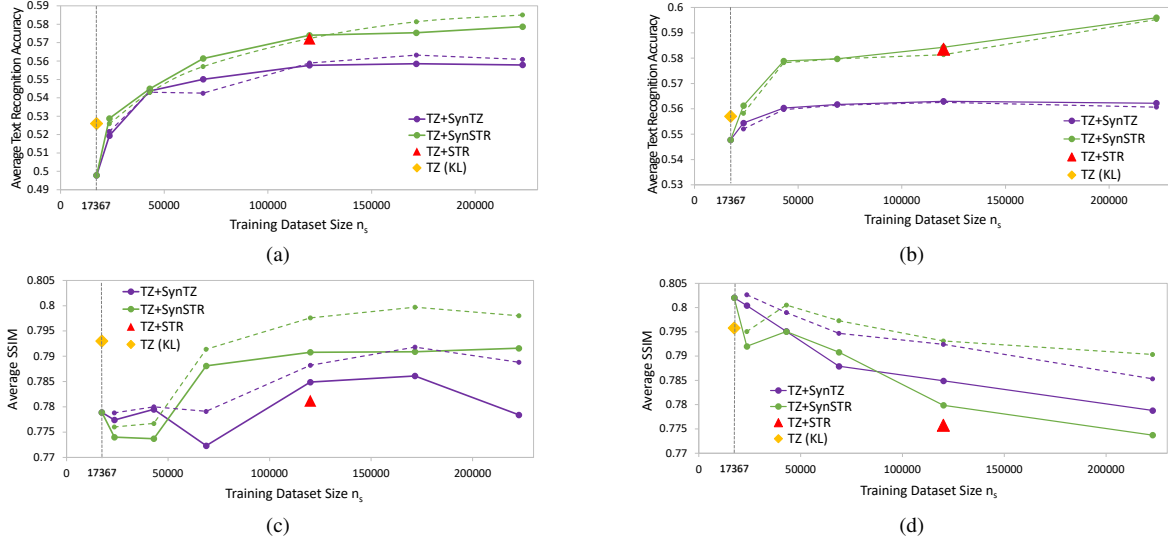


Figure 3. Evaluation results of TATT and the text-conditional DM trained on the augmented datasets. (a) and (b) show average recognition accuracy of TATT and the text-conditional DM, respectively. (c) and (d) show average SSIM of TATT and the text-conditional DM, respectively. Dotted lines show the results of fine-tuning using TextZoom only. The solid and dotted lines in the same color correspond to the same augmented dataset. CTC loss was used to train the text prior generator except in the case of “TZ (KL).” Here, “TZ (KL)” indicates the case where KL loss was used. The size of TextZoom is 17,367.



Figure 4. Examples of artistic, curved, or rotated text images excluded by preprocessing.

dictionary. A comparison based on various maximum word lengths is available in Sec. D. Words comprised solely of digits were included with a 10% probability. The character set consists of digits as well as lowercase and uppercase letters. To fine-tune the text prior generator, CTC loss [13] was used. Hereafter, recognition accuracy is evaluated by CRNN and presented as an average of the results for the three difficulty levels of the TextZoom test set. All other experimental conditions and evaluation metrics were the same as those in Sec. 4.1.

Datasets. When training Super-resolver and Degradar, we used LR-HR paired text images of TextZoom. For training Synthesizer, two datasets were prepared: the first is the HR images of TextZoom, and the second is the HR images of TextZoom plus datasets for scene text recognition. The scene text recognition datasets consist of 11 real labeled datasets, including SVT [50], IIIT [31], IC13 [21], IC15 [20], COCO [49], RCTW [44], Uber [58], ArT [7], LSVT [47], MLT19 [33], and ReCTS [57]. We used the above datasets with preprocessing conducted by the authors of [1] to remove irregular text images (e.g., those containing non-English characters or vertical text). The preprocessed scene text recognition dataset contains 276K text images.

Preprocessing for Synthesizer. The preprocessed dataset introduced earlier contains some text images that are difficult to handle in the proposed framework, as shown in Fig. 4. To remove these text images, we used CRNN [42] pretrained with synthetic datasets [14, 18], which has poor recognition ability for non-simple text images, as shown in Fig. 4. We removed text images whose text labels did not match those predicted by CRNN. As a result, the preprocessed dataset contains 103K text images.

Postprocessing for Synthesized Text Images. In text image synthesis, there were some instances where the input texts did not align with the texts in the synthesized images. Therefore, we conducted postprocessing similar to preprocessing. Specifically, ASTER [43] was used to predict the texts of the synthesized text images. If the predicted texts did not match the input texts, the synthesized image was removed. We believe that ASTER is well-suited for this postprocessing because it can achieve state-of-the-art performance and has no strong language model to correct incorrect spellings.

4.2.1 Effectiveness of Dataset Augmentation

For the preparation of text images, we explored three methods. (1) using the preprocessed scene text recognition dataset introduced earlier. (2) using Synthesizer trained exclusively on TextZoom. (3) using Synthesizer trained on both TextZoom and the preprocessed scene text recognition dataset. We named these three datasets STR, SynTZ,

SynSTR, respectively. The resulting text images served as input to Super-resolver and Degradar, generating corresponding paired images. These paired images were then combined with the original TextZoom (TZ) to form augmented datasets. Consequently, the first augmented dataset is termed TZ+STR, the second as TZ+SynTZ, and the third as TZ+SynSTR. We denote the size of the augmented dataset as n_s and that of TextZoom as n_t .

TATT was first used to evaluate the augmented dataset. Figure 3a shows the average recognition accuracy versus n_s as solid lines. Recognition accuracy tends to improve as n_s increases. In addition, the recognition accuracy when trained on TZ+SynSTR was higher than when trained on TZ+SynTZ. TZ+STR achieved an accuracy similar to TZ+SynSTR with the same dataset size. Additionally, Fig. 3c shows SSIM versus n_s . SSIM did not improve depending on n_s . Moreover, SSIM is worse overall compared with the case where $n_s = n_t$ and KL loss was used. However, this can be improved by fine-tuning with TextZoom only. Fine-tuning results are shown as dotted lines. We can see that fine-tuning is effective when the augmented dataset size is large. The results of PSNR are provided in Sec. C.

The text-conditional DM was next used to evaluate the augmented dataset. Figures 3b and 3d show the average recognition accuracy and SSIM versus n_s , respectively. Similarly to the results of TATT, recognition accuracy tends to improve as n_s increases. However, SSIM tends to decrease as n_s increases. This decrease of SSIM can be attributed the difference that HR images of TextZoom contain many blurred images while those of the augmented dataset do not. We consider that this is not a negative effect because removing blurred images can improve recognition accuracy. Additional experimental results on this trade-off are provided in Sec. F. Similarly to the results of TATT, the decrease of SSIM can be reduced by fine-tuning with TextZoom only.

5. Ablation Studies

5.1. Comparison to Synthetic Degradation Methods

Synthetic LR images can be obtained by applying simple degradation techniques, such as bicubic interpolation and blur kernels. In this subsection, we compare the paired text images generated by these simple degradation techniques with those generated by the proposed framework. We employed synthetic degradation pipelines presented in BSRGAN [56] and Real-ESRGAN [53] to generate synthetic LR-HR paired images online and trained TATT using these images. The evaluation involved three datasets: SynTZ, SynSTR, and STR, introduced in Sec. 4.2.1. Note that the original TextZoom was not included in these datasets to ensure a fair comparison. The comparison results are presented in Tab. 3. As can be seen, the utilization of paired

Method	SSIM ($\times 10^{-2}$)			PSNR		
	SynTZ	SynSTR	STR	SynTZ	SynSTR	STR
BSRGAN	73.05	75.50	74.56	21.42	21.87	21.36
Real-ESRGAN	72.29	74.29	72.47	21.15	21.20	20.12
Ours	73.95	76.39	75.97	20.40	21.15	20.66

(a)

Method	Acc. (%)		
	SynTZ	SynSTR	STR
BSRGAN	40.82	46.50	47.62
Real-ESRGAN	40.36	44.50	45.03
Ours	48.46	52.69	54.03

(b)

Table 3. Comparison with synthetic degradation methods. TATT was trained with these paired images, and the performance was evaluated in terms of average of (a) SSIM/PSNR and (b) recognition accuracy on TextZoom.

images generated by our framework results in significantly higher recognition accuracy compared to those generated by the BSRGAN and Real-ESRGAN pipelines. On the other hand, their SSIM and PSNR scores remain competitive.

5.2. KL Loss vs. CTC Loss

KL loss [28] and CTC loss [13] were used to train the text prior generator using TextZoom and the augmented dataset, respectively. This is because the appropriate dataset size varies depending on the loss function. Figure 6 compares the recognition accuracy with the KL and CTC losses for different training dataset sizes. We can see that the original TextZoom is not sufficiently large when CTC loss is used while CTC loss is more effective than KL loss for a larger size of the augmented dataset. The CTC loss tends to induce overfitting when the dataset size is small because it directly measures the distance from the ground-truth texts. On the other hand, KL loss measures the distance from the probability predicted for HR images and is less likely to induce overfitting while the distance may not be accurate.

5.3. Effects of Super-resolver and Degradar

Synthesizer was trained on HR text images of TextZoom and the preprocessed scene text recognition dataset. These datasets contain blurred images as well as clear images, resulting in Synthesizer generating both blurred and clear images. Therefore, without Super-resolver and Degradar, it is anticipated that the proposed framework cannot create stable-quality paired images. Here, we refer to the images generated by Synthesizer as medium-resolution (MR) images. To examine the effectiveness of the LR-HR paired images, we created two extended datasets: one without Super-resolver (LR-MR paired images) and the other without Degradar (MR-HR paired images). We trained TATT on these datasets and compared them with the results of LR-HR paired images. As shown in Tab. 4, the highest recognition accuracy was achieved when LR-HR paired images were used. Although the highest SSIM/PSNR were

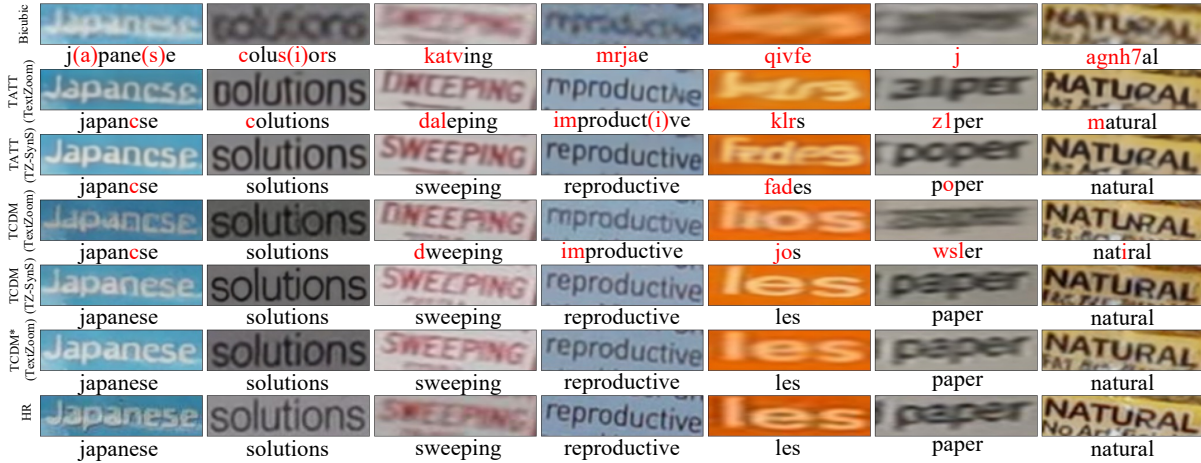


Figure 5. Examples of SR text images generated from TATT and text-conditional DMs. The training datasets used are shown in parentheses. Red characters indicate incorrect or missing results.




Figure 6. Comparison between KL and CTC losses. It shows the average recognition accuracy of TATT trained with KL and CTC losses for varying training dataset sizes.

Input-Target	SSIM ($\times 10^{-2}$)	PSNR	Acc. (%)
LR-MR	79.08	22.26	54.40
MR-HR	78.32	22.00	52.74
LR-HR	78.49	21.92	55.77

Table 4. Comparison of LR-HR pairs with two different input-target pairs: LR-MR and MR-HR pairs. TATT was trained on each of the three pairs and the performance was evaluated on the test set of TextZoom. MR images represent images generated by Synthesizer trained on TextZoom only.

obtained with the LR-MR pairs, we consider that this is due to the domain gap between the synthesized and the original HR images (see Sec. 4.2.1 and F for more details).

6. Limitation

Super-resolver and Degradator must be trained on TextZoom. Thus, non-simple text images, which are not included in TextZoom such as those shown in Fig. 4, cannot be handled in the proposed framework. To alleviate this drawback, extensions such as combining with image degradation methods that do not require paired images can be




Figure 7. Examples of text images generated by the proposed framework. The last two columns show the LR-HR paired images and the first column shows the images generated by Synthesizer.

considered.

7. Conclusions

We experimentally showed that text-conditional DMs are effective in STISR, achieving state-of-the-art results in the TextZoom evaluation. Notably, the performance of the text-conditional DMs becomes remarkably enhanced when trained using ground-truth texts. Leveraging this exceptionally expressive capacity, we proposed a novel framework for synthesizing LR-HR paired text images. Our proposed framework, encompassing three distinct text-conditional DMs, can generate high-quality paired text images from user-provided texts. Our experiments demonstrated a marked improvement in the performance of STISR methods when trained using the paired images generated through our proposed framework.

References

- [1] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3113–3122, 2021. **6**
- [2] Dmitry Baranchuk, Andrey Voynov, Ivan Rubachev, Valentin Khulkov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. In *Proceedings of the International Conference on Learning Representations*, 2022. **3**
- [3] Darwin Bautista and Rowel Atienza. Scene text recognition with permuted autoregressive sequence models. In *Proceedings of the European Conference on Computer Vision*, page 178–196, 2022. **2**
- [4] Ahmed Cheikh Sidiya, Xuan Xu, Ning Xu, and Xin Li. Degradation learning and skip-transformer for blind face restoration. *Frontiers in Signal Processing*, 3, 2023. **2**
- [5] Jingye Chen, Bin Li, and Xiangyang Xue. Scene text telescope: Text-focused scene image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12026–12035, June 2021. **2, 5**
- [6] Jingye Chen, Haiyang Yu, Jianqi Ma, Bin Li, and Xiangyang Xue. Text gestalt: Stroke-aware scene text image super-resolution. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 285–293, 2022. **2, 5**
- [7] Chee Kheng Chng, Yuliang Liu, Yipeng Sun, Chun Chet Ng, Canjie Luo, Zihan Ni, ChuanMing Fang, Shuaitao Zhang, Junyu Han, Errui Ding, et al. Icdar2019 robust reading challenge on arbitrary-shaped text-rrc-art. In *Proceedings of the IEEE International Conference on Document Analysis and Recognition*, pages 1571–1576, 2019. **6**
- [8] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *arXiv preprint arXiv:2209.04747*, 2022. **12**
- [9] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *Advances in Neural Information Processing Systems*, volume 34, 2021. **3, 13**
- [10] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. **1, 2, 5**
- [11] Chao Dong, Ximei Zhu, Yubin Deng, Chen Change Loy, and Yu Qiao. Boosting optical character recognition: A super-resolution approach. *arXiv preprint arXiv:1506.02211*, 2015. **2**
- [12] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7094–7103, 2021. **2**
- [13] Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the International Conference on Machine Learning*, page 369–376, 2006. **2, 6, 7**
- [14] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2315–2324, 2016. **6**
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. **1, 3, 12**
- [16] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23:47–1, 2022. **3**
- [17] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS Workshop on Deep Generative Models and Downstream Applications*, 2021. **3**
- [18] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. In *NeurIPS Workshop on Deep Learning*, 2014. **6**
- [19] Yuming Jiang, Shuai Yang, Haonan Qju, Wayne Wu, Chen Change Loy, and Ziwei Liu. Text2human: Text-driven controllable human image generation. *ACM Transactions on Graphics*, 41(4), 2022. **3**
- [20] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, Faisal Shafait, Seiichi Uchida, and Ernest Valveny. Icdar competition on robust reading. In *Proceedings of 13th International Conference on Document Analysis and Recognition*, pages 1156–1160, 2015. **6**
- [21] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazàn Almazàn, and Lluís Pere de las Heras. Icdar competition on robust reading. In *Proceedings of 12th International Conference on Document Analysis and Recognition*, pages 1484–1493, 2013. **6**
- [22] 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**
- [23] Jiwon Kim, Jung Kwon Lee, and Kyoung Mu Lee. Accurate image super-resolution using very deep convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1646–1654, 2016. **1**
- [24] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 624–632, 2017. **1**
- [25] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE*

- Conference on Computer Vision and Pattern Recognition*, July 2017. 1, 2, 5
- [26] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. 1
- [27] Canjie Luo, Lianwen Jin, and Zenghui Sun. Moran: A multi-object rectified attention network for scene text recognition. *Pattern Recognition*, 90:109–118, 2019. 2, 5
- [28] Jianqi Ma, Shi Guo, and Lei Zhang. Text prior guided scene text image super-resolution. *arXiv preprint arXiv:2106.15368*, 2021. 2, 3, 5, 7, 12
- [29] Jianqi Ma, Zhetong Liang, and Lei Zhang. A text attention network for spatial deformation robust scene text image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5911–5920, June 2022. 2, 3, 5
- [30] Shunta Maeda. Unpaired image super-resolution using pseudo-supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 291–300, 2020. 2
- [31] Anand Mishra, Alahari Karteek, and C. V. Jawahar. Scene text recognition using higher order language priors. In *Proceedings of the British Machine Vision Conference*, 2009. 6
- [32] Yongqiang Mou, Lei Tan, Hui Yang, Jingying Chen, Leyuan Liu, Rui Yan, and Yaohong Huang. Plugnet: Degradation aware scene text recognition supervised by a pluggable super-resolution unit. In *Proceedings of the European Conference on Computer Vision*, pages 158–174, 2020. 2
- [33] Nibal Nayef, Yash Patel, Michal Busta, Pinaki Nath Chowdhury, Dimosthenis Karatzas, Wafa Khlif, Jiri Matas, Uma-pada Pal, Jean-Christophe Burie, Cheng-lin Liu, et al. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition—rrc-mlt-2019. In *Proceedings of the IEEE International conference on document analysis and recognition*, pages 1582–1587, 2019. 6
- [34] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. In *Proceedings of the International Conference on Machine Learning*, volume 162, 2022. 1, 3
- [35] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *Proceedings of the International Conference on Machine Learning*, volume 139, pages 8162–8171, 2021. 3, 5, 13
- [36] Yuhui Quan, Jieting Yang, Yixin Chen, Yong Xu, and Hui Ji. Collaborative deep learning for super-resolving blurry text images. *IEEE Transactions on Computational Imaging*, 6:778–790, 2020. 2
- [37] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 1, 3
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, June 2022. 4, 12
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the Medical Image Computing and Computer-Assisted Intervention*, pages 234–241, 2015. 3
- [40] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. 1, 3, 4, 12
- [41] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–14, 2022. 3
- [42] 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 Transactions on Pattern Analysis and Machine Intelligence*, 39(11):2298–2304, 2017. 2, 5, 6
- [43] Baoguang Shi, Mingkun Yang, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Aster: An attentional scene text recognizer with flexible rectification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(9):2035–2048, 2019. 2, 5, 6
- [44] Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. Icdar2017 competition on reading chinese text in the wild (rctw-17). In *Proceedings of the IEEE international conference on document analysis and recognition*, volume 1, pages 1429–1434, 2017. 6
- [45] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the International Conference on Machine Learning*, volume 37, pages 2256–2265, 2015. 1, 3, 12
- [46] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in Neural Information Processing Systems*, 32, 2019. 1
- [47] Yipeng Sun, Zihan Ni, Chee-Kheng Chng, Yuliang Liu, Canjie Luo, Chun Chet Ng, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, et al. Icdar 2019 competition on large-scale street view text with partial labeling-rrc-lsvt. In *Proceedings of the IEEE International Conference on Document Analysis and Recognition*, pages 1557–1562, 2019. 6
- [48] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. 4
- [49] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016. 6
- [50] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *Proceedings of the IEEE Inter-*

- national Conference on Computer Vision*, pages 1457–1464, 2011. [6](#)
- [51] Wenjia Wang, Enze Xie, Xuebo Liu, Wenhai Wang, Ding Liang, Chunhua Shen, and Xiang Bai. Scene text image super-resolution in the wild. In *Proceedings of the European Conference on Computer Vision*, pages 650–666, 2020. [1](#), [2](#), [5](#)
- [52] Wenjia Wang, Enze Xie, Peize Sun, Wenhai Wang, Lixun Tian, Chunhua Shen, and Ping Luo. Textsr: Content-aware text super-resolution guided by recognition. *CoRR*, abs/1909.07113, 2019. [2](#)
- [53] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 1905–1914, 2021. [2](#), [7](#)
- [54] Xiangyu Xu, Deqing Sun, Jinshan Pan, Yujin Zhang, Hanspeter Pfister, and Ming-Hsuan Yang. Learning to super-resolve blurry face and text images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 251–260, 2017. [2](#)
- [55] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Yingxia Shao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *arXiv preprint arXiv:2209.00796*, 2022. [12](#)
- [56] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4791–4800, 2021. [2](#), [7](#)
- [57] Rui Zhang, Yongsheng Zhou, Qianyi Jiang, Qi Song, Nan Li, Kai Zhou, Lei Wang, Dong Wang, Minghui Liao, Mingkun Yang, et al. Icdar 2019 robust reading challenge on reading chinese text on signboard. In *Proceedings of the IEEE International conference on document analysis and recognition*, pages 1577–1581, 2019. [6](#)
- [58] Ying Zhang, Lionel Gueguen, Ilya Zharkov, Peter Zhang, Keith Seifert, and Ben Kadlec. Uber-text: A large-scale dataset for optical character recognition from street-level imagery. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, volume 2017, page 5, 2017. [6](#)
- [59] Yongbing Zhang, Siyuan Liu, Chao Dong, Xinfeng Zhang, and Yuan Yuan. Multiple cycle-in-cycle generative adversarial networks for unsupervised image super-resolution. *IEEE transactions on Image Processing*, 29:1101–1112, 2019. [2](#)
- [60] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 2018. [1](#)
- [61] 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 *Proceedings of the ACM International Conference on Multimedia*, page 2908–2917, 2021. [2](#), [5](#)
- [62] Minyi Zhao, Miao Wang, Fan Bai, Bingjia Li, Jie Wang, and Shuigeng Zhou. C3-stir: Scene text image super-resolution with triple clues. In *Proceedings of the International Joint Conferences on Artificial Intelligence*, pages 1707–1713, 2022. [2](#), [3](#), [5](#)
- [63] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2223–2232, 2017. [2](#)