

# Recognition-Synergistic Scene Text Editing

Zhengyao Fang<sup>\*1</sup>, Pengyuan Lyu<sup>\*3</sup>, Jingjing Wu<sup>\*4</sup>, Chengquan Zhang<sup>3</sup>, Jun Yu<sup>1</sup>,  
Guangming Lu<sup>1</sup>, Wenjie Pei<sup>†1,2</sup>

<sup>1</sup>Harbin Institute of Technology, Shenzhen

<sup>2</sup>Peng Cheng Laboratory

<sup>3</sup>Tencent

<sup>4</sup>Department of Computer Vision Technology, Baidu Inc.

{zhengyaonineve, jingjingwu.hit, wenjiecoder}@outlook.com

lvpyuan@gmail.com, zchengquan@gmail.com, {yujun, luguangm}@hit.edu.cn

## Abstract

Scene text editing aims to modify text content within scene images while maintaining style consistency. Traditional methods achieve this by explicitly disentangling style and content from the source image and then fusing the style with the target content, while ensuring content consistency using a pre-trained recognition model. Despite notable progress, these methods suffer from complex pipelines, leading to suboptimal performance in complex scenarios. In this work, we introduce Recognition-Synergistic Scene Text Editing (**RS-STE**), a novel approach that fully exploits the intrinsic synergy of text recognition for editing. Our model seamlessly integrates text recognition with text editing within a unified framework, and leverages the recognition model’s ability to implicitly disentangle style and content while ensuring content consistency. Specifically, our approach employs a multi-modal parallel decoder based on transformer architecture, which predicts both text content and stylized images in parallel. Additionally, our cyclic self-supervised fine-tuning strategy enables effective training on unpaired real-world data without ground truth, enhancing style and content consistency through a twice-cyclic generation process. Built on a relatively simple architecture, **RS-STE** achieves state-of-the-art performance on both synthetic and real-world benchmarks, and further demonstrates the effectiveness of leveraging the generated hard cases to boost the performance of downstream recognition tasks. Code is available at <https://github.com/ZhengyaoFang/RS-STE>.

## 1. Introduction

Scene Text Editing (STE) aims to modify the textual content in scene text images while preserving the original style.

<sup>\*</sup>Equal contribution.

<sup>†</sup>Corresponding author.

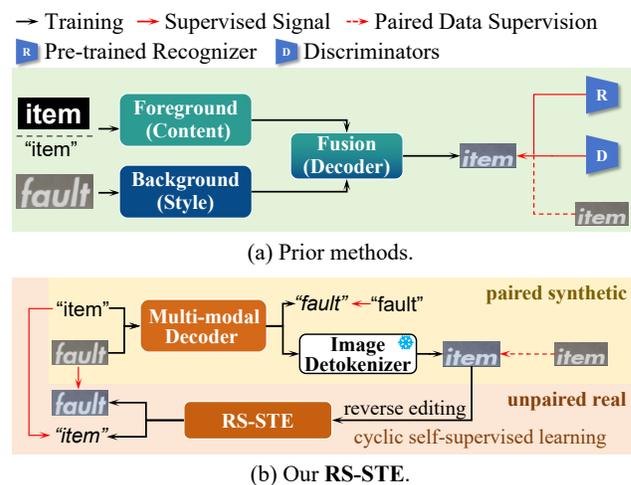


Figure 1. Prior methods for scene text editing involve intricate modeling for explicit separation of text content and background style. In contrast, our **RS-STE** conducts synergistic modeling of scene text recognition and text editing in a unified framework, which allows for implicit text-style separation while ensuring content consistency. Besides, the specially designed Cyclic Self-supervised Fine-tuning enables effective training of **RS-STE** on unpaired real-world data, substantially enhancing the generalizability in real-world scenarios.

This technology holds significant potential, enabling designers to efficiently edit and replace textual information in images. Additionally, it can be applied to image generation, thereby enhancing the performance of other Optical Character Recognition (OCR) tasks, such as text detection and recognition. Given its importance, STE has garnered increasing attention from researchers.

STE presents two primary challenges. Firstly, the diverse appearances of scene text, including variations in background, font, and layout, pose significant difficulties for STE. Secondly, the lack of paired real training data neces-

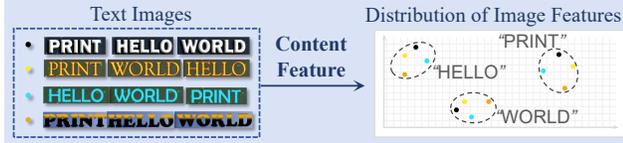


Figure 2. Distribution of some content features extracted by our **RS-STE**. Images with the same text content but different background styles become closer in the encoded feature space of a recognition model, implying the capability of recognition models to separate style from content.

sitates training existing methods on synthetic data. The domain gap between synthetic and real data hinders the ability of models trained on synthetic data to generalize effectively to real-world scenarios. Many methods have been proposed to address the aforementioned challenges. Most of these methods follow a pipeline that involves explicitly decomposing the information in the source image into style and content components, and then merging the source style with the target text content to produce the target image. A pre-trained recognition model is always used to ensure the content consistency in the edited image. For instance, as illustrated in Figure 1 (a), existing approaches [8, 31, 34, 42, 45] typically involve an explicit separation of the background and foreground images. Other methods [24, 47] focus on learning to disentangle the content features and style features from the source style image. Furthermore, to address the issue of lacking real paired data, several learning mechanisms [23, 24, 31, 44] have been proposed. These mechanisms employ a combination of real and synthetic data during training to enhance the model’s generalization capabilities in real-world scenarios. By leveraging both types of data, these approaches aim to bridge the domain gap and improve the performance of scene text editing models when applied to real-world images. Most of these methods, when training on unpaired real data, follow a paradigm similar to the illustration in Figure 1 (a), where paired data supervision is not available.

While previous methods have made remarkable progress, their intricate pipelines introduce two potential limitations that hinder further performance improvements. Firstly, explicitly separating style and content is a challenging task and may not always be perfectly accurate, which can result in suboptimal outcomes when these components are recombined. Secondly, these methods often consist of multiple interconnected modules, and the joint optimization of these modules can also lead to less-than-ideal results.

In this paper, we introduce **RS-STE**, a novel approach that not only addresses the limitations of existing methods but also delivers superior editing outcomes. The impetus for our approach arises from a fundamental observation: recognition models inherently separate style from content, as il-

lustrated in Figure 2. By capitalizing on this characteristic, we seamlessly integrate the recognition model with the editing model within a cohesive framework. Specifically, we have developed a multi-modal parallel decoder based on the transformer decoder architecture. This decoder, upon receiving the encoded tokens of the specified text and the source style image, concurrently predicts the text content on the style image and generates an image with the source style and specified text. Furthermore, we propose a Cyclic Self-Supervised Fine-tuning mechanism, as illustrated in Figure 1 (b) for *unpaired real* data, to effectively harness real-world data for training purposes. This design maximizes the potential of the recognition model in two principal ways: (1) It implicitly decouples style and content, allowing the model to better capture attributes for generating target image, and (2) Within the Cyclic Self-Supervised Fine-tuning, the recognition model’s supervision ensures the consistency of the generated content. As a result, our approach offers a significant advantage: it eliminates the need for multiple modules to explicitly decouple style and content, and it obviates the necessity for a separate recognition model to verify the accuracy of the generated content. This greatly simplifies the overall pipeline and circumvents the challenges faced by previous methodologies.

To conclude, our contributions are listed as follows:

- We propose a simple yet effective scene text editing method dubbed **RS-STE**, which conducts recognition-synergistic scene text editing in a unified framework. Such design enables implicit separation between background style and text content, thereby eliminating intricate model design.
- We design the Cyclic Self-supervised Fine-tuning Strategy, which allows for effective training on unpaired real-world data to substantially enhance its generalizability in real-world scenarios.
- Our **RS-STE** achieves state-of-the-art performance on both synthetic and real-world scene text editing benchmarks. We further validated the effectiveness of our generated image on downstream recognition tasks.

## 2. Related Work

### 2.1. Scene Text Editing

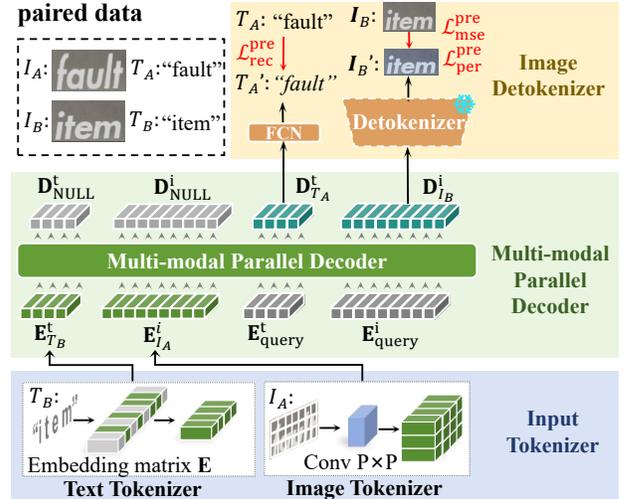
Scene Text Editing refers to the process of modifying text within natural scene images while maintaining the visual consistency and context of the surrounding elements. Early work designed complex modules to explicitly separate foreground and background. SRNet [42] leverages three respective modules for learning background reconstruction, render text, and final fusion. Similarly, STEFANN [34] focuses on character-wise rendering through a text conversion and color transfer module, applying them to the inpainted background. Advancing these foundations, Swap-

Text [45] introduces a spatial transformation to adapt to oriented text. To implicitly decouple the content and style features and use the unpaired real-world data for training, TextStyleBrush [23] propose to use task-adaptive StyleGAN2 [21] along with self-supervised training strategy. RewriteNet [24] learns to separate content and style features and fine-tunes on real-world images with a self-supervised training scheme. MOSTEL [31] intentionally focuses on style by discarding content information, using style augmentation techniques to merge style and content images, thus producing target content with a reconstructed background. CLASTE [44] uses extra background restoration module for background restoration and integrates the background with foreground content. DARLING [47] combines aligned style and content features using a multi-task decoder enhanced by self-attention blocks, showcasing a blend of synthesis, self-supervised learning, and contextual awareness in modern scene text editing techniques. Recently, STEEM [8] introduces a minimized background reconstruction technique to explicitly decouple the style and content and further enhance text editing fidelity.

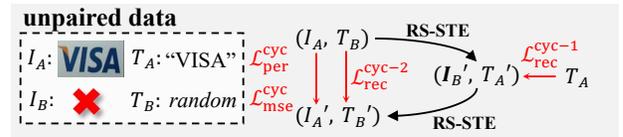
Recently, several works [4, 5, 16, 40] explore the manipulation of text within the entire image using stable diffusion [33]. They encompass the intricate aspects of layout arrangement and the accurate rendering of textual content. As this focus differs from the concerns addressed in our paper, we have not conducted a comparison with our work.

## 2.2. MLLM for Image Generation and Editing

In response to the notable progress of large language models in natural language processing [1, 2, 6], the field of multi-modal large language models (MLLM) has made significant strides in recent years. MLLMs leverage both natural language and visual inputs, allowing these models to understand and manipulate visual data guided by textual descriptions. This dual-modality capability builds upon foundational image generation models, such as GANs [12] and diffusion models [15], but advances them by incorporating language as a critical component in model design. Recent works [1, 11, 25, 37, 39] have developed architectures capable of processing text and image modalities simultaneously, achieving a more nuanced integration of linguistic and visual information. These approaches demonstrate enhanced performance in image generation tasks, where MLLMs generate high-quality visuals that align closely with the semantic content of textual prompts. Furthermore, some MLLMs [3, 13, 28] offer innovative capabilities for image editing by enabling users to adjust existing images through descriptive language, such as modifying attributes or inserting new elements, rather than relying on pixel-level manipulation. Inspired by these methods, our approach integrates the multi-modal language model **RS-STE**, which is specialized in scene text editing.



(a) Fully-Supervised Pre-training Stage of **RS-STE** on Paired Datasets



(b) Self-Supervised Cyclic Fine-tuning Stage on Unpaired Datasets

Figure 3. (a) illustrates the model structure of **RS-STE** and the fully-supervised pre-training stage using paired synthetic datasets. (b) depicts the cyclic self-supervised fine-tuning stage with unpaired real-world datasets.

## 3. Method

### 3.1. Overview

The aim of scene text editing is to edit text image  $I_A$  to synthesize image  $I_B$  by altering the text content  $T_A$  into the desired content  $T_B$  while retaining the image style of  $I_A$ . Our proposed **RS-STE** for this task is able to conduct text recognition and editing within a unified framework, resulting in a straightforward pipeline. As shown in Figure 3, it comprises Input Tokenizer, Multi-modal Parallel Decoder, and Image Detokenizer.

Given the target text content  $T_B$  and a reference image  $I_A$ , Input Tokenizer encodes them into text embeddings and image embeddings respectively, and outputs a cascaded embedding sequence. Then Multi-modal Parallel Decoder performs scene text editing in the feature space and predicts the tokens of  $T_A'$  and  $I_B'$  in a parallel manner. Lastly, Image Detokenizer generates target image  $I_B'$  from decoded image tokens  $D_{I_B'}^i$ . While the generated  $I_B'$  contains different text content from  $I_A$ , their stylistic components including background and typeface are required to be completely identical.

Our **RS-STE** is optimized in two learning stages. It is first trained on a large corpus of synthesized data with paired  $I_A$  and  $I_B$  to endow it with the basic capability of

scene text editing. Then in the second stage, it is further optimized on unpaired real-world data (without ground-truth) using our specially designed cyclic self-supervised fine-tuning strategy, which substantially improves its robustness and generalizability towards real-world data. We will first describe the model structure of **RS-STE** in Section 3.2, and then elaborate on the training strategy in Section 3.3.

### 3.2. RS-STE

**Input Tokenizer.** The Input Tokenizer encodes the input target text  $T_B$  and the reference style image  $I_A$  separately. For text encoding, we learn an embedding matrix  $\mathbf{E} \in \mathbb{R}^{(|\Sigma|+1) \times C}$  for alphabet  $\Sigma$ , from which we can encode  $T_B = \{c_1, \dots, c_L\}$  by indexing the corresponding character embeddings sequentially, thereby obtaining the text embedding  $\mathbf{E}_{T_B}^t \in \mathbb{R}^{L \times C}$ .

We adopt the ViT-based tokenization approach to encode the reference style image  $I_A \in \mathbb{R}^{H \times W \times 3}$ . To be specific, we apply a convolutional layer with a kernel size of  $P \times P$  to split image into  $\frac{H}{P} \times \frac{W}{P}$  patches and capture visual information, producing flattened visual feature sequence  $\mathbf{E}_{I_A}^i \in \mathbb{R}^{N \times C}$ , where  $N = (HW)/P^2$ .

**Multi-modal Parallel Decoder (MMPD).** A scene text recognition model is capable of extracting text-related features from an image by implicitly distinguishing between text and background style. In light of this, instead of disentangling style and text content via separate task modeling as other methods [8, 31, 34, 42, 45] perform, our **RS-STE** model conducts both scene text recognition and text editing in the unified Multi-modal Parallel Decoder to leverage the synergy of text recognition towards editing. As shown in Figure 3, given  $\mathbf{E}_{T_B}^t$  and  $\mathbf{E}_{I_A}^i$ , the Multi-modal Parallel Decoder is optimized to recognize the text content  $T'_A$  while performing text editing in the feature space to predict the token features of the target image  $I'_B$ .

The Multi-modal Parallel Decoder is designed in the structure of Transformer decoder. Following the classical modeling paradigm of multi-modal language foundation models [7, 11, 43], we initialize learnable query embeddings corresponding to the text and image prediction, denoted as  $\mathbf{E}_{\text{query}}^t \in \mathbb{R}^{L \times C}$  and  $\mathbf{E}_{\text{query}}^i \in \mathbb{R}^{N \times C}$  respectively. They are sequentially concatenated after  $\mathbf{E}_{T_B}^t$  and  $\mathbf{E}_{I_A}^i$ , and fed into the Multi-modal Parallel Decoder:

$$\begin{aligned} & [\mathbf{D}_{\text{NULL}}^t, \mathbf{D}_{\text{NULL}}^i, \mathbf{D}_{T_A}^t, \mathbf{D}_{I_B}^i] \\ & = \mathcal{F}_{\text{MMPD}}([\mathbf{E}_{T_B}^t, \mathbf{E}_{I_A}^i, \mathbf{E}_{\text{query}}^t, \mathbf{E}_{\text{query}}^i]). \end{aligned} \quad (1)$$

where  $\mathbf{D}_{T_A}^t \in \mathbb{R}^{L \times C}$  and  $\mathbf{D}_{I_B}^i \in \mathbb{R}^{N \times C}$  are the decoded token features for  $T_A$  and  $I_B$ , respectively. It is noteworthy that the first  $(L + N)$  output tokens aligned with  $\mathbf{E}_{T_B}^t$  and  $\mathbf{E}_{I_A}^i$ , denoted as  $\mathbf{D}_{\text{NULL}}^t$  and  $\mathbf{D}_{\text{NULL}}^i$ , are not used since they cannot access the full content of  $\mathbf{E}_{T_B}^t$  and  $\mathbf{E}_{I_A}^i$ .  $\mathbf{D}_{T_A}^t$  is further used to perform text recognition by a fully connected

layer (FCN) that predicts the character probabilities  $\mathbf{P}_{T_A} \in \mathbb{R}^{L \times (|\Sigma|+1)}$ , while  $\mathbf{D}_{I_B}^i$  is fed into the Image Detokenizer to synthesize the edited image  $I_B$ .

**Image Detokenizer.** We utilize the pre-trained VAE decoder of LDM [33] as the Image Detokenizer and fine-tune it on the synthesized training data. Following the routine training paradigm [9], the Image Detokenizer is fine-tuned before the training of the Input Tokenizer and Multi-modal Parallel Decoder of **RS-STE** for stable optimization.

### 3.3. Training Strategy

Our **RS-STE** is optimized in two stages: a fully-supervised pre-training stage on paired synthetic data and a cyclic self-supervised fine-tuning stage on unpaired real-world data.

**Fully-Supervised Pre-training Stage.** As collecting paired real-world data for supervised scene text editing is infeasible, we first pre-train our **RS-STE** on synthetic paired data to equip it with the fundamental capability of scene text editing. Since our model is able to perform synergistic modeling of both scene text recognition and text editing, we conduct supervised learning on both tasks, as illustrated in Figure 3 (a). To be specific, we adopt cross-entropy loss to optimize scene text recognition:

$$\mathcal{L}_{\text{rec}}(T_A, T'_A) = -\frac{1}{L} \sum_{i=1}^L \sum_{c=1}^{|\Sigma|+1} G(T_A)[i, c] \log(\mathbf{P}_{T_A}[i, c]), \quad (2)$$

where  $L$  is the max length of text.  $G(T_A)[i, c]$  is the one-hot encoding ground truth at position  $i$ , with the  $c$ -th character in the pre-defined alphabet equal to 1, while  $\mathbf{P}_{T_A}[i, c]$  is the corresponding prediction by our **RS-STE**.

To supervise scene text editing, we employ the mean squared error (MSE) loss for pixel-level supervision and perceptual loss [18] for semantic alignment between the edited image and the ground truth. Formally, for the edited image  $I'_B$ , the MSE loss and perceptual loss are defined as:

$$\begin{aligned} \mathcal{L}_{\text{mse}}(I_B, I'_B) &= \|I_B - I'_B\|_2^2, \\ \mathcal{L}_{\text{per}}(I_B, I'_B) &= \mathbb{E} [\|\phi_i(I_B) - \phi_i(I'_B)\|_2^2], \end{aligned} \quad (3)$$

where  $\phi_i$  is the features extracted from *relu1\_2*, *relu2\_2*, *relu3\_3* and *relu4\_3* of a pre-trained VGG-16 network [36].

Integrating all three losses, the overall learning objective for our **RS-STE** in the pre-training stage is formulated as:

$$\mathcal{L}^{\text{pre}} = \lambda_1 \mathcal{L}_{\text{rec}}^{\text{pre}}(T_A, T'_A) + \lambda_2 \mathcal{L}_{\text{mse}}^{\text{pre}}(I_B, I'_B) + \lambda_3 \mathcal{L}_{\text{per}}^{\text{pre}}(I_B, I'_B), \quad (4)$$

where  $\lambda_1$ ,  $\lambda_2$  and  $\lambda_3$  are balancing weights, and the superscript ‘pre’ indicates that these losses only apply to the pre-training stage.

**Cyclic Self-Supervised Fine-tuning Stage.** Despite the abundance of paired synthetic training data available for

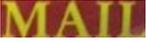
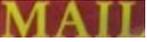
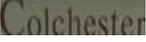
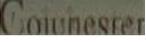
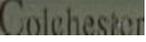
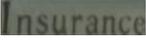
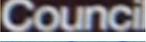
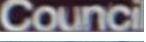
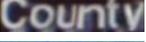
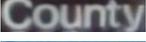
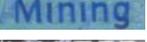
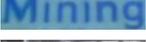
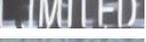
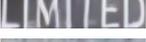
source	target	SRNet	MOSTEL	TextCtrl	RS-STE
	“ROYAL”				
	“Insurance”				
	“County”				
	“Mining”				
	“LIMITED”				
	“FESTIVE”				
	“ORWELL”				
	“group”				

Figure 4. Editing examples compared with other methods.

pre-training, the significant disparity between synthetic data and real-world data severely limits the applicability of the pre-trained model in real-world scenarios. Nevertheless, conducting direct supervised learning with real-world data is impractical in the absence of paired data for scene text editing. To circumvent this problem, we devise the Cyclic Self-Supervised Fine-tuning strategy, which conducts scene text editing twice on a same style image with reverse operation by interchanging the target text, reproducing the initial style image. As shown in Figure 3 (b), given a style image  $I_A$  and the target text  $T_B$ , our **RS-STE** generates  $I'_B$  and predicts  $T'_A$  in the first scene text editing. Then, using  $I'_B$  and  $T'_A$  as the style image and target text respectively, **RS-STE** performs the second editing and obtains the edited image  $I'_A$  and recognized text  $T'_A$ , which should be the reproduction of the initial style image  $I_A$ . The whole process can be expressed as:

$$\begin{aligned} (I'_B, T'_A) &= \mathcal{F}_{\text{RS-STE}}(I_A, T_B), \\ (I'_A, T'_B) &= \mathcal{F}_{\text{RS-STE}}(I'_B, T'_A), \end{aligned} \quad (5)$$

where  $\mathcal{F}_{\text{RS-STE}}$  denotes editing function by our **RS-STE**. The proposed cyclic editing procedure allows us to perform supervision on the reproduced image  $I'_A$ , which is equivalent to self-supervised learning.

In this stage, we also use MSE loss and perceptual loss to supervise the generation of  $I'_A$ . Meanwhile, we apply the recognition loss to both predicted  $T'_A$  and  $T'_B$  in twice text editing, which can prevent the model from collapsing into an identical mapping before and after cyclic editing:

$$\begin{aligned} \mathcal{L}^{\text{cyc}} &= \lambda_4 \mathcal{L}_{\text{mse}}^{\text{cyc}}(I_A, I'_A) + \lambda_5 \mathcal{L}_{\text{per}}^{\text{cyc}}(I_A, I'_A) \\ &+ \lambda_6 \mathcal{L}_{\text{rec}}^{\text{cyc-1}}(T_A, T'_A) + \lambda_7 \mathcal{L}_{\text{rec}}^{\text{cyc-2}}(T_B, T'_B), \end{aligned} \quad (6)$$

where  $\lambda_4$ ,  $\lambda_5$ ,  $\lambda_6$  and  $\lambda_7$  are hyper-parameters for balancing between different losses.

## 4. Experiment

### 4.1. Datasets

**Training.** In the pre-training stage, we utilize the open-source synthetic Tamper-train [31] dataset, which consists of 150k images. Additionally, following [47], we incorporate 4M paired synthetic data samples generated using the same image synthesis engine<sup>1</sup> employed in Tamper-train. During the fine-tuning stage, we utilize the MLT-2017 [29] real dataset in accordance with the MOSTEL [31] to ensure a fair comparison. In order to further explore the effectiveness of our method, we also conducted training on the Union14M-L [17] dataset to verify whether more complex and diverse training data can yield better results.

**Evaluation.** For evaluation, we use the paired synthetic dataset Tamper-Syn2k [31] and the paired real-world dataset ScenePair [46] to assess the discrepancies between our edited images and their target counterparts. These two datasets contain 2,000 pairs and 1,280 pairs of text images, respectively, with different text content but the same stylistic components. Additionally, we leverage the unpaired real-world dataset Tamper-Scene [31], comprising 7,725 unpaired images, to indirectly evaluate editing quality. Furthermore, we incorporate six commonly used text recognition benchmark datasets—IIIT 5K-Words (IIIT) [27], ICDAR2013 (IC13) [19], Street View Text (SVT) [41], ICDAR2015 (IC15) [20], Street View Text-Perspective (SVTP) [30], and CUTE80 (CUTE) [32]—for further indirect evaluation of editing quality in more complex and diverse real-world scenarios. Finally, we utilize the Union-benchmark [17] dataset to assess improvements in recognition models facilitated by our targeted data augmentation strategy as detailed in Section 4.6.

<sup>1</sup><https://github.com/youdao-ai/SRNet-Datagen>

Table 1. Comparison on editing performance with state-of-the-art methods on paired synthetic dataset Tamper-Syn2k, unpaired real dataset Tamper-Scene and paired real dataset ScenePair. The SSIM and SeqAcc are presented in percent (%).

Methods	Tamper-Syn2k				Tamper-Scene	ScenePair					#Params
	MSE↓	PSNR↑	SSIM↑	FID↓	RecAcc↑	MSE↓	PSNR↑	SSIM↑	FID↓	RecAcc↑	
pix2pix	0.1450	9.18	34.15	127.21	13.26	—	—	—	—	—	54.4
SRNet[42]	0.0216	18.66	49.97	64.37	30.26	0.0561	14.08	26.66	49.22	17.84	<b>18.95</b>
SwapText[45]	0.0194	19.43	52.43	—	54.83	—	—	—	—	—	—
MOSTEL[31]	0.0135	20.27	56.94	33.79	66.54	0.0519	14.46	27.45	49.19	37.69	54.0
DARLING[47]	0.0120	20.80	60.07	44.48	70.85	—	—	—	—	—	23.7
STEEM [48]	0.0122	20.83	72.10	<b>24.67</b>	78.80	—	—	—	—	—	69.6
TextCtrl[46]	0.0130	20.79	66.60	31.13	74.17	0.0447	14.99	37.56	43.78	84.67	1216.0
Ours	<b>0.0076</b>	<b>22.54</b>	<b>72.90</b>	30.29	<b>86.12</b>	<b>0.0267</b>	<b>17.35</b>	<b>46.09</b>	<b>41.37</b>	<b>91.80</b>	54.4

Table 2. **Text Recognition Accuracy** on the STR common benchmark datasets. *Base* is provided as baseline results of the recognition model [10] on the original images, implying the upper bound of recognition performance. Others are the recognition results on the edited datasets generated by different models.

Methods	Real Training	Recognition Benchmarks Accuracy						Avg
	Dataset	IIT	IC13	SVT	IC15	SVTP	CUTE	
Base	—	96.5	95.1	94.7	85.9	89.5	89.6	91.8
TextCtrl	—	70.0	68.8	73.7	62.6	63.4	58.7	66.2
MOSTEL	MLT2017	48.2	40.6	41.1	34.4	21.6	35.1	36.8
Ours (w/o $\mathcal{L}^{acc}$ )	—	61.2	61.2	66.8	53.8	46.4	44.8	55.7
Ours	MLT2017	76.6	74.8	89.6	83.0	86.4	<b>80.6</b>	81.8
Ours	Union14M-L	<b>78.4</b>	<b>76.2</b>	<b>91.3</b>	<b>85.8</b>	<b>88.2</b>	77.4	<b>82.9</b>

## 4.2. Implementation Details

We initialize the pre-trained VAE [22] with configuration parameters  $f = 4$ ,  $Z = 8192$  and  $d = 3$  from LDM<sup>2</sup>. This model is fine-tuned on the Tamper-train dataset, and its decoder is frozen for use in subsequent stages. The minGPT model<sup>3</sup>, with 22.5M parameters, serves as the foundation of MMPD and is pre-trained from scratch. We adopt AdamW [26] as the optimizer with  $\beta_1 = 0.9$  and  $\beta_2 = 0.95$ . Following [47], the image size in both training and evaluation is set to  $32 \times 128$ .

In the fully-supervised pre-training stage, we set a batch size of 32 and a learning rate of  $1.44 \times 10^{-4}$ , training for a total of 300k iterations.  $\lambda_1$ ,  $\lambda_2$ , and  $\lambda_3$  are set to 1, 10, and 1, respectively. In the subsequent cyclic self-supervised fine-tuning stage, we use a batch size of 16 and a learning rate of  $7.2 \times 10^{-5}$ , training for 1k iterations. The weights  $\lambda_4$ ,  $\lambda_5$ ,  $\lambda_6$  and  $\lambda_7$  are set to 10, 1, 50, 50, respectively. All experiments are conducted on 4 NVIDIA 3090 GPUs.

## 4.3. Evaluation Metrics

Regarding the paired evaluation datasets Tamper-Syn2k and ScenePair, where the target edited images are specified, we utilize four metrics to assess the differences between the predicted and target images. These metrics include 1) Mean

Squared Error (MSE), which measures the  $L_2$  distance; 2) Peak Signal-to-Noise Ratio (PSNR), representing the ratio of peak signal power to noise power; 3) Structural Similarity Index Measure (SSIM), which evaluates mean structural similarity; and 4) Fréchet Inception Distance [14] (FID), which calculates the distance between features extracted using the InceptionV3 [38] model. Higher PSNR, SSIM, and lower MSE, FID indicate better performance.

Moreover, for the real datasets Tamper-Scene and ScenePair, we use the same pre-trained recognition model CRNN [35], as in MOSTEL [40] to compute the Recognition Accuracy (RecAcc) of the generated images, which serves as an indirect indicator of the consistency between the text content of the edited images and their target texts.

## 4.4. Comparison on Editing Performance

We perform two sets of experiments to evaluate our method: 1) evaluation on the standard benchmarks for scene text editing including paired synthetic dataset Tamper-Syn2k, paired real dataset ScenePair, and unpaired real dataset Tamper-Scene; 2) experiments on more challenging real datasets, i.e., the classical datasets for text recognition.

Table 1 presents the scene text editing performance of different methods in the first set of experiments. Our method achieves the best performance across all evaluation metrics, except for the FID on synthetic data. Specifically, on the paired real dataset ScenePair, **RS-STE** shows significant improvements in MSE, PSNR, SSIM, and RecAcc. On the unpaired dataset Tamper-Scene, we observe a 7.32% increase in RecAcc compared to the state-of-the-art method STEEM [48]. For the synthetic dataset Tamper-Syn2k, since some of the ground truth images are not even discernible to the naked eye, our method fails to generate images that closely resemble the ground truth. This leads to a visual discrepancy and results in a higher FID.

In the second set of evaluations, we assess our method on common benchmarks for text recognition in terms of RecAcc metric, which is highly challenging for scene text editing. Using images in these datasets as the style images, we randomly sample a word that differs from the image annota-

<sup>2</sup><https://github.com/CompVis/latent-diffusion>

<sup>3</sup><https://github.com/karpathy/minGPT>

tion to serve as the target text for each style image. We then apply various models to perform scene text editing. Finally, we employ a more powerful text recognition model ABINet [10] to conduct text recognition on the edited images, as CRNN [35] shows limited performance on these more challenging datasets. Table 2 presents the experimental results. Note that the recognition performance of ABINet [10] on the original images, labeled as *Base*, is used as baseline implying the upper bound of recognition performance. We observe that our **RS-STE** demonstrates an average performance improvement of 45.0% over MOSTEL [31], utilizing the same real training data MLT2017. It is encouraging that, when fine-tuned on Union14M-L dataset, our **RS-STE** achieves comparable performance with the upper bound on ‘SVT’, ‘IC15’ and ‘SVTP’, which further demonstrates the effectiveness of our method.

To further illustrate the effectiveness of our method, we also provide qualitative comparisons. Accordingly, we present several visual qualitative examples in Figure 4. It is clear that our method significantly outperforms other methods in terms of editing effects. Further visualization results can be found in the supplementary material.

#### 4.5. Ablation Study

**Intrinsic Recognition.** To validate the effectiveness of **RS-STE**’s intrinsic recognition, which implicitly disentangles style and content while simultaneously ensuring content accuracy, we conduct a series of comparative experiments. The results are presented in Table 3 and 4.

First, to assess the disentanglement ability of our **RS-STE**, we conduct an experiment under the same conditions optimized for editing task, disregarding the model’s recognition ability. Specifically, we achieve this by excluding the recognition loss during the training process. The comparison results are listed in Table 3 *w/o*  $\mathcal{L}_{rec}^{pre}$  and *Ours*. As can be easily seen, joint optimization of recognition and editing abilities during training can significantly enhance the editing performance of the model, particularly in terms of overall structure and realism, resulting in a 3.20 increase in SSIM and a 3.67 decrease in FID.

Second, to verify that the intrinsic recognition model offers superior content consistency compared to a conventional pre-trained external recognition model, we conduct an experiment incorporating the pre-trained model ABINet [10] solely to supervise the recognition of generated images. Results under identical training conditions are presented in Table 3 *w/*  $\mathcal{L}_{rec}^{per}$ . It can be observed that compared to the lack of supervised recognition task *w/o*  $\mathcal{L}_{rec}^{pre}$ , using an external recognition model for supervision can improve model performance. However, the external recognition model can only constrain the consistency of the edited results in terms of content and does not achieve the decoupling between style and content like our intrinsic recogni-

Table 3. Ablation studies on  $\mathcal{L}^{pre}$  applied in the fully-supervised pre-training stage using paired synthetic Tamper-Syn2k dataset. ‘*w/o*  $\mathcal{L}_{mse}^{pre}$ ’ denotes the absence of  $\mathcal{L}_{mse}^{pre}$ , while ‘*w/*  $\mathcal{L}_{rec}^{pre}$ ’ means the utilization of pre-trained recognition supervisor.

Methods	Tamper-Syn2k			
	MSE↓	PSNR↑	SSIM↑	FID↓
<i>w/o</i> $\mathcal{L}_{per}^{pre}$	<b>0.0058</b>	<b>24.05</b>	<b>78.02</b>	70.06
<i>w/o</i> $\mathcal{L}_{mse}^{pre}$	0.0086	21.94	73.79	<b>29.63</b>
<i>w/o</i> $\mathcal{L}_{rec}^{pre}$	0.0082	22.26	69.70	33.96
<i>w/</i> $\mathcal{L}_{rec}^{pre}$	0.0079	22.44	70.71	31.73
Ours	0.0076	22.54	72.90	30.29

Table 4. Ablation studies on  $\mathcal{L}^{cyc}$  applied in the cyclic self-supervised fine-tuning stage using unpaired real dataset Tamper-Scene and paired real dataset ScenePair.

Methods	Tamper-Scene	ScenePair				
	RecAcc↑	MSE↓	PSNR↑	SSIM↑	FID↓	RecAcc↑
<i>w/o</i> $\mathcal{L}_{per}^{cyc}$	<b>96.23</b>	0.0431	15.16	15.75	134.31	88.52
<i>w/o</i> $\mathcal{L}_{mse}^{cyc}$	83.79	0.0422	14.79	17.64	97.00	69.53
<i>w/o</i> $\mathcal{L}_{rec}^{cyc-1}$	0.00	0.0500	14.51	22.10	27.79	4.38
<i>w/o</i> $\mathcal{L}_{rec}^{cyc-2}$	0.00	0.0492	14.76	22.10	24.61	4.38
<i>w/o</i> $\mathcal{L}^{cyc}$	69.01	0.0309	16.66	36.81	46.60	80.63
Ours	86.12	<b>0.0267</b>	<b>17.35</b>	<b>46.09</b>	<b>41.37</b>	<b>91.80</b>

tion model, thus resulting in slightly inferior performance.

**Cyclic Self-Supervised Fine-tuning Strategy.** Cyclic training enables **RS-STE** to optimize in a self-supervised manner on unpaired real-world data, thereby significantly enhancing its editing performance on real data. Specifically, as shown in Table 2 and 4 *w/o*  $\mathcal{L}^{cyc}$ , without fine-tuning on real datasets, **RS-STE** can only exhibit an average editing recognition accuracy of 55.7% on commonly used text recognition benchmarks, 69.01% on Tamper-Scene, and poor style consistency on ScenePair dataset. This is primarily due to the significant domain gap between the synthetic dataset used for pre-training and real-world scenarios. However, as shown in Table 2, when fine-tuned on the MLT2017 dataset using a cyclic training strategy, our approach achieves an average editing recognition accuracy of 81.8%. Moreover, fine-tuning with the more complex and diverse Union14M-L dataset improves accuracy to 82.9%, and significantly enhances style consistency on real-world dataset ScenePair, as shown in Table 4 *Ours*. These results highlight the considerable potential for performance gains in our method.

Additionally, during the cyclic training stage, our intrinsic recognition ensures content consistency between the target text and the edited image. As shown in Table 4 and Figure 5, when our model does not utilize  $\mathcal{L}_{rec}^{cyc-1}$  and  $\mathcal{L}_{rec}^{cyc-2}$  for constraints, it tends to learn an identical mapping from the original image to itself, which directly results in the loss



Figure 5. Visualization examples of ablation study.

Table 5. Performance improvements of classical recognition models yielded from fine-tuning with edited bad cases from scene text editing models as data augmentation. All methods are pre-trained on Union14M-L.

Methods	Augmentation Model	Union14M-Benchmark							Avg.
		Curve	Multi-Oriented	Artistic	Contextless	Salient	Multi-Words	General	
ABINet [10]	–	73.0	51.0	64.6	72.7	70.4	61.6	77.9	67.3
	MOSTEL [31]	73.7 +0.7	53.1 +2.1	65.0 +0.4	<b>73.8 +1.1</b>	72.2 +1.8	60.1 -1.5	78.0 +0.1	68.0 +0.7
	Ours	<b>74.5 +1.5</b>	<b>54.5 +3.5</b>	<b>65.8 +1.2</b>	73.7 +1.0	<b>73.9 +3.5</b>	<b>65.4 +3.8</b>	<b>78.8 +0.9</b>	<b>69.5 +2.2</b>
MAERec-S [17]	–	81.4	71.4	72.0	82.0	78.5	82.4	82.5	78.6
	MOSTEL [31]	84.0 +2.6	72.0 +0.6	72.9 +0.9	80.2 -1.8	79.1 +0.6	82.3 -0.1	82.1 -0.4	78.9 +0.3
	Ours	<b>85.0 +4.6</b>	<b>75.4 +4.0</b>	<b>76.1 +4.1</b>	<b>82.9 +0.9</b>	<b>80.9 +2.4</b>	<b>84.3 +1.9</b>	<b>83.2 +0.7</b>	<b>81.1 +2.5</b>

of our model’s ability to perform scene text editing. This demonstrates that, through the intrinsic recognition supervision in the cyclic training stage, our model is capable of decomposing content and style on real-world data, while ensuring content consistency.

**Loss Functions.** In addition to recognition loss, MSE and Perception losses also play important roles for preserving style. We also conduct ablation studies to specifically discuss the effectiveness of different losses in the pre-training and cyclic training stages. The results are listed in Table 3 and Table 4 respectively.

As shown in Table 3, in the pre-training stage,  $w/o \mathcal{L}_{per}^{pre}$  or  $w/o \mathcal{L}_{mse}^{pre}$  will significantly result in poorer editing performance. Specifically,  $\mathcal{L}_{per}^{pre}$  enhances the visual realism of generated images, as indicated by lower FID scores, while  $\mathcal{L}_{mse}^{pre}$  ensures the pixel-level similarity indicated by MSE, PSNR, and SSIM. The same phenomenon can be seen in Figure 5 (a).

In the cyclic fine-tuning stage, cyclic perceptual loss ( $\mathcal{L}_{per}^{cyc}$ ) and cyclic MSE loss ( $\mathcal{L}_{mse}^{cyc}$ ) supervise the pixels of the reproduced results. Without these losses, as shown in Table 4  $w/o \mathcal{L}_{per}^{cyc}$  and  $w/o \mathcal{L}_{mse}^{cyc}$ , though **RS-STE** can achieve a better recognition accuracy on Tamper-Scene, it fails to preserve the style and tend to generate the targeted content in an extraneous style, as illustrated in Figure 5 (b)  $w/o \mathcal{L}_{per}^{cyc}$ .

#### 4.6. Targeted Data Augmentation for Recognition

In this section, we address the generation of targeted training data that simulates challenging cases encountered by the recognition model. This data augmentation strategy aims to fine-tune the recognition model, thereby increasing its accuracy and robustness in real-world applications.

By addressing specific recognition errors, this targeted fine-tuning strategy markedly improves both general recognition models, such as ABINet [10], and state-of-the-art recognition models, including MAERec-S [17]. The results in Table 5 show a significant improvement: the average recognition accuracy of ABINet [10] and MAERec-S [17] increase by 2.2% and 2.5%, respectively, with our augmentation method. In contrast, MOSTEL [31] only leads to an improvement of 0.7% and 0.3%, respectively. This result illustrates that our targeted data augmentation technique using our method significantly enhances the performance of recognition, even when the recognition model already achieves strong results. Implementation details will be provided in the supplementary.

## 5. Conclusion

In this work, we present **RS-STE**, which conducts recognition-synergistic scene text editing in a unified framework, thereby eliminating the intricate model design for explicit disentanglement of background style and text content. Moreover, we devise the Cyclic Self-Supervised Fine-Tuning strategy, which is able to fine-tune our **RS-STE** using unpaired real-world data, significantly enhancing its generalizability to real-world scenarios. Extensive experiments validate the effectiveness of the proposed **RS-STE**.

## Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (62372133, 62125201, U24B20174), in part by Shenzhen Fundamental Research Program (Grant NO. JCYJ20220818102415032).

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3
- [2] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 3
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, pages 18392–18402, 2023. 3
- [4] Haoxing Chen, Zhuoer Xu, Zhangxuan Gu, Yaohui Li, Changhua Meng, Huijia Zhu, Weiqiang Wang, et al. Diffute: Universal text editing diffusion model. *NIPS*, 36, 2024. 3
- [5] Jingye Chen, Yupan Huang, Tengchao Lv, Lei Cui, Qifeng Chen, and Furu Wei. Textdiffuser: Diffusion models as text painters. *NIPS*, 36, 2024. 3
- [6] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *JMLR*, 24(240):1–113, 2023. 3
- [7] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *arXiv preprint arXiv:2309.11499*, 2023. 4
- [8] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, Xiangwen Kong, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Dreamllm: Synergistic multimodal comprehension and creation. In *ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 2, 3, 4
- [9] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, pages 12873–12883, 2021. 4
- [10] 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. 6, 7, 8
- [11] Yuying Ge, Sijie Zhao, Jinguo Zhu, Yixiao Ge, Kun Yi, Lin Song, Chen Li, Xiaohan Ding, and Ying Shan. Seed-x: Multimodal models with unified multi-granularity comprehension and generation. *arXiv preprint arXiv:2404.14396*, 2024. 3, 4
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NIPS*, 27, 2014. 3
- [13] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 3
- [14] 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. *NIPS*, 30, 2017. 6
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NIPS*, 33:6840–6851, 2020. 3
- [16] Jiabao Ji, Guanhua Zhang, Zhaowen Wang, Bairu Hou, Zhifei Zhang, Brian Price, and Shiyu Chang. Improving diffusion models for scene text editing with dual encoders. *arXiv preprint arXiv:2304.05568*, 2023. 3
- [17] Qing-Yuan Jiang, Jiapeng Wang, Dezhi Peng, Chongyu Liu, and Lianwen Jin. Revisiting scene text recognition: A data perspective. *ICCV*, pages 20486–20497, 2023. 5, 8
- [18] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711. Springer, 2016. 4
- [19] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. In *2013 12th international conference on document analysis and recognition*, pages 1484–1493. IEEE, 2013. 5
- [20] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. In *2015 13th international conference on document analysis and recognition (ICDAR)*, pages 1156–1160. IEEE, 2015. 5
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, pages 8110–8119, 2020. 3
- [22] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. 6
- [23] Praveen Krishnan, Rama Kovvuri, Guan Pang, Boris Vasilev, and Tal Hassner. Textstylebrush: transfer of text aesthetics from a single example. *PAMI*, 45(7):9122–9134, 2023. 2, 3
- [24] Junyeop Lee, Yoonsik Kim, Seonghyeon Kim, Moonbin Yim, Seung Shin, Gayoung Lee, and Sungrae Park. Rewritenet: Reliable scene text editing with implicit decomposition of text contents and styles. *arXiv preprint arXiv:2107.11041*, 2021. 2, 3
- [25] Yanwei Li, Yuechen Zhang, Chengyao Wang, Zhisheng Zhong, Yixin Chen, Ruihang Chu, Shaoteng Liu, and Jiaya Jia. Mini-gemini: Mining the potential of multi-modality vision language models. *arXiv preprint arXiv:2403.18814*, 2024. 3
- [26] I Loshchilov. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [27] Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. In *BMVC-British machine vision conference*. BMVA, 2012. 5
- [28] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images

- using guided diffusion models. In *CVPR*, pages 6038–6047, 2023. 3
- [29] Nibal Nayef, Fei Yin, Imen Bizid, Hyunsoo Choi, Yuan Feng, Dimosthenis Karatzas, Zhenbo Luo, Umapada Pal, Christophe Rigaud, Joseph Chazalon, et al. Icdar2017 robust reading challenge on multi-lingual scene text detection and script identification-rrc-mlt. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, pages 1454–1459. IEEE, 2017. 5
- [30] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. In *ICCV*, pages 569–576, 2013. 5
- [31] Yadong Qu, Qingfeng Tan, Hongtao Xie, Jianjun Xu, Yuxin Wang, and Yongdong Zhang. Exploring stroke-level modifications for scene text editing. In *AAAI*, pages 2119–2127, 2023. 2, 3, 4, 5, 6, 7, 8
- [32] Anhar Risnumawan, Palaiahnakote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications*, 41(18):8027–8048, 2014. 5
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10684–10695, 2022. 3, 4
- [34] Prasun Roy, Saumik Bhattacharya, Subhankar Ghosh, and Umapada Pal. Stefann: scene text editor using font adaptive neural network. In *CVPR*, pages 13228–13237, 2020. 2, 4
- [35] 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. *PAMI*, 39(11):2298–2304, 2016. 6, 7
- [36] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 4
- [37] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiyang Yu, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. In *CVPR*, pages 14398–14409, 2024. 3
- [38] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *CVPR*, pages 2818–2826, 2016. 6
- [39] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 3
- [40] Yuxiang Tuo, Wangmeng Xiang, Jun-Yan He, Yifeng Geng, and Xuansong Xie. Anytext: Multilingual visual text generation and editing. *arXiv preprint arXiv:2311.03054*, 2023. 3, 6
- [41] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. In *ICCV*, pages 1457–1464. IEEE, 2011. 5
- [42] Liang Wu, Chengquan Zhang, Jiaming Liu, Junyu Han, Jingtuo Liu, Errui Ding, and Xiang Bai. Editing text in the wild. In *ACMMM*, pages 1500–1508, 2019. 2, 4, 6
- [43] Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*, 2023. 4
- [44] Fuxiang Yang, Tonghua Su, Xiang Zhou, Donglin Di, Zhongjie Wang, and Songze Li. Self-supervised cross-language scene text editing. In *ACMMM*, pages 4546–4554, 2023. 2, 3
- [45] Qiangpeng Yang, Jun Huang, and Wei Lin. Swaptxt: Image based texts transfer in scenes. In *CVPR*, pages 14700–14709, 2020. 2, 3, 4, 6
- [46] Weichao Zeng, Yan Shu, Zhenhang Li, Dongbao Yang, and Yu Zhou. Textctrl: Diffusion-based scene text editing with prior guidance control. *arXiv preprint arXiv:2410.10133*, 2024. 5, 6
- [47] Boqiang Zhang, Hongtao Xie, Zuan Gao, and Yuxin Wang. Choose what you need: Disentangled representation learning for scene text recognition removal and editing. In *CVPR*, pages 28358–28368, 2024. 2, 3, 5, 6
- [48] Jianqun Zhou, Pengwen Dai, Yang Li, Manjiang Hu, and Xiaochun Cao. Explicitly-decoupled text transfer with the minimized background reconstruction for scene text editing. *TIP*, pages 1–1, 2024. 6