

The Surprisingly Straightforward Scene Text Removal Method With Gated Attention and Region of Interest Generation: A Comprehensive Prominent Model Analysis

Hyeonsu Lee¹[0000–0002–6317–9883] and Chankyu Choi¹[0000–0002–9166–2100]

NAVER Corp {hyeon-su.lee, chankyu.choi}@navercorp.com

Abstract. Scene text removal (STR), a task of erasing text from natural scene images, has recently attracted attention as an important component of editing text or concealing private information such as ID, telephone, and license plate numbers. While there are a variety of different methods for STR actively being researched, it is difficult to evaluate superiority because previously proposed methods do not use the same standardized training/evaluation dataset. We use the same standardized training/testing dataset to evaluate the performance of several previous methods after standardized re-implementation. We also introduce a simple yet extremely effective Gated Attention (GA) and Region-of-Interest Generation (RoIG) methodology in this paper. GA uses attention to focus on the text stroke as well as the textures and colors of the surrounding regions to remove text from the input image much more precisely. RoIG is applied to focus on only the region with text instead of the entire image to train the model more efficiently. Experimental results on the benchmark dataset show that our method significantly outperforms existing state-of-the-art methods in almost all metrics with remarkably higher-quality results. Furthermore, because our model does not generate a text stroke mask explicitly, there is no need for additional refinement steps or sub-models, making our model extremely fast with fewer parameters. The dataset and code are available at <https://github.com/naver/garnet>.

1 Introduction

Scene text removal (STR) is a task of erasing text from natural scene images, which is useful for privacy protection, text editing in images/videos, and Augmented Reality (AR) translation.

A lot of current STR research utilizes the deep learning method. Early studies [22] attempted to erase all text in an image without using a text region mask, but this produced imprecise, dissatisfactory results such as some regions being blurred out. Furthermore, it was impossible to selectively erase letters in a certain region with this method, which drastically limited applications. In order to overcome these limitations, recent studies [17, 21] in the field built text removal




Fig. 1. Visualization of the proposed model’s input and output. Visualization of the results of applying attention to the text stroke and applying attention to the text stroke surrounding regions found by the GA. We obtained satisfying STR results using RoIG.

models after generating text region masks through manual or automatic means. This lead to comparably better quality results and wider applications.

However, these past studies did not use the same standardized training dataset and evaluation dataset, making it impossible to evaluate performance in a fair manner. For instance, some papers trained on different subsets of synthetic data[5]. Furthermore, they do not have the same input image dimensions, which affects speed, accuracy, and model parameters. Thus, it is difficult to select a previously proposed model for the unique requirements of a specific application when there is no standardized comparison of the model size and speed.

In this paper, we perform a fair comparison with both qualitative (text removal quality) and quantitative (model size, inference time) metrics by re-implementing prominent previously proposed models with the same standards, training them with the same training dataset, and evaluating them on the same evaluation dataset. Furthermore, we introduce the *Gated Attention* (GA) and the *Region of Interest Generation* (RoIG). GA is a simple yet extremely effective method that uses attention to focus on the text stroke and its surroundings, while RoIG drastically improves the training efficiency of the STR model.

Figure 1 shows GA found the Text Stroke Region (TSR) and the Text Stroke Surrounding Region (TSSR) after training the pseudo mask and how RoIG generates the results only for text mask regions. More details are described in Section 4.

Our contributions can be summarized as follows:

1. We re-implemented previously proposed methods with the same input image size, trained them on the same training dataset, and compared their performance with the same evaluation dataset. We also performed comparisons of speed and model size.
2. We proposed a method called GA which not only distinguishes between the background and the text stroke, but also utilizes attention to identify both

Table 1. A side-by-side comparison of how our proposed model differs from previous works.

Method	Use text box	Selective removal	Stage	Stroke localization	Surround localization	RoI Generation
EnsNet [22]	x	x	1	x	x	x
MTRNet [18]	o	o	1	x	x	x
MTRNet++ [17]	o	o	2	o	x	x
EraseNet [9]	x	x	2	x	x	x
Tang <i>et al.</i> [15]	o	o	2	o	x	o(crop)
Ours	o	o	1	o	o	o

TSR and TSSR. To the best of our knowledge, there is no other previous study that considered both TSR and TSSR while performing STR.

3. We proposed a method called RoIG, which helps the network focus on only the region with text instead of the entire image to train the model more efficiently.
4. The proposed method generates higher quality results than existing state-of-the-art methods on both synthetic and real datasets. It is also significantly lighter and faster than most other prominent previously proposed popular methods because our model does not have additional refinement steps or sub-models needed to generate a text stroke mask.

2 Related Works

The major trend in scene text removal before the emergence of deep learning was traditional rule-based methods [2, 16], which are often hand-crafted and require prior domain knowledge.

Recently, deep learning-based text removal has been proposed by adopting popular GAN-based methods. EnsNet [22] is simple and fast because it doesn’t need any auxiliary inputs. However, its results are blurry and of low quality. Furthermore, its practical applications are limited because it is impossible to only erase text in a certain region. MTRNet [18] requires the generation of text box region masks through manual or automatic means. These studies show that text region masks can improve the network’s performance but cannot guarantee high-quality results. MTRNet++ [17] and EraseNet [9] proposed a coarse-refinement two-stage network. While the results are of higher quality, the model is much bigger, slower, and too complicated.

Zdenek *et al.* [21] used an auxiliary text detector to retrieve the text box mask, then attempted to erase text through a general inpainting method. However, they were unable to generate results of satisfactory quality because they did not consider qualities specific to text like the text strokes. Tang *et al.* [15] can train their model in a rather efficient manner because they erase text by cropping only the text regions. However, this method ignores all global contexts other than the cropped region and has difficulty precisely cropping the region of curved texts.

Table 2. A comparison of our model’s performance with previous STR models on real data (SCUT-EnsText [9]). We compare the results claimed in the original papers of the previous prominent STR models with the results we re-measured in an equal environment. We realize that there are fundamental differences in models and that it would be unfair to evaluate them under potentially disadvantageous circumstances. To make things truly fair, we also pasted the text region (produced by the model using a text box mask) over the original image, then evaluated the results yet again. The values right of “/” reflect experiments conducted after this adjustment was made. Our model produces superior results in all metrics regardless. The best score is highlighted in bold.

	Method	Train Data	Size	Image-Eval		
				PSNR	SSIM	AGE
Reported	Scene Text Eraser [12]	Real	256	25.47/-	90.14/-	6.01/-
	EnsNet [22]	Real	512	29.53/-	92.74/-	4.16/-
	MTRNet [18]	Syn(75%)	256	-/-	-/-	-/-
	MTRNet++ [17]	Syn(95%)	256	-/-	-/-	-/-
	EraseNet [9]	Real	512	32.30/37.26	95.42/96.86	3.02/-
Our experiment	EnsNet [22]	Real+Syn	512	31.05/32.99	94.78/95.16	2.67/1.85
	MTRNet [18]	Real+Syn	256	30.61/36.06	89.85/95.72	3.92/1.21
		Real+Syn	512	32.46/36.89	95.86/96.41	3.12/0.97
	MTRNet++ [17]	Real+Syn	256	35.29/37.40	96.31/96.68	1.26/1.09
		Real+Syn	512	34.86/36.50	96.32/96.51	1.48/1.35
	EraseNet [9]	Real+Syn	512	30.54/37.16	96.27/97.53	3.07/1.20
	EraseNet [9] + M	Real+Syn	512	34.29/40.18	97.73/97.98	2.28/0.69
	Ours	Real+Syn	512	-/41.37	-/98.46	-/0.64

We designed a model that can focus on only the text region area while also managing to take the global context of the entire image into consideration. This model does not require an additional refinement process or a sub-model dedicated to text stroke localization, leading to a drastically faster and lighter model.

3 Comprehensive Prominent Model Analysis

In this section, we analyze the difference between previous methods and ours. We also show the results of evaluating previous methods with one standardized dataset.

Table 1 outlines the specific differences between our method and previously proposed methods. First, our proposed method takes a text box mask as the input to our model, leading to results of significantly better quality as well as the option for users to selectively erase only the text that they wish to. Second, our proposed method can localize the TSR and TSSR to erase text in a surgical manner. Finally, the use of RoIG makes our proposed method show off significantly better results than previous methods without even implementing a coarse-refinement 2-stage network.

Table 3. A comparison of our model’s performance with previous STR models on real data (SCUT-EnsText [9]). R, P, and F refer to recall, precision and F-score, respectively. The best score is highlighted in bold.

	Method	Train Data	Size	Detection-Eval			GPU Time(ms)	Params
				P	R	F		
Reported	Scene Text Eraser [12]	Real	256	40.9	5.9	10.2		
	EnsNet [22]	Real	512	68.7	32.8	44.4	-	12.4M
	MTRNet [18]	Syn(75%)	256	-	-	-	-	54.4M
	MTRNet++ [17]	Syn(95%)	256	-	-	-	-	18.7M
	EraseNet [9]	Real	512	53.2	4.6	8.5	-	-
Our experiment	EnsNet [22]	Real+Syn	512	73.1	54.7	62.6	12.0	12.4M
	MTRNet [18]	Real+Syn	256				21.9	50.3M
		Real+Syn	512	69.8	41.1	51.2	51.3	
	MTRNet++ [17]	Real+Syn	256				51.3	
		Real+Syn	512	58.6	20.5	30.4	238.7	
	EraseNet [9]	Real+Syn	512	40.8	6.3	10.9	47.4	17.8M
	EraseNet [9] + M	Real+Syn	512	37.3	6.1	10.3	47.4	17.8M
	Ours	Real+Syn	512	15.5	1.0	1.8	14.9	12.4M

Table 2, Table 3 and Table 4 show the performance of several previous STR methods on real and synthetic data. Details of the experiment and re-implementation are in Section 5 and Section 5.3.

After performing an objectively fair comparison, we found that our method generates higher quality results than existing state-of-the-art methods on both synthetic and real datasets. It is also significantly lighter and faster than any other method except EnsNet [22].

Table 4. A comparison of our model’s performance with previous STR models on synthetic data (Oxford [5]). The notation is the same as Table 2 and Table 3.

Method	Input Size	Image-Eval			Detection-Eval		
		PSNR	SSIM	AGE	P	R	F
Original Image	512				71.3	51.5	59.8
EnsNet [22]	512	36.67/39.74	97.71/97.94	1.25/0.77	55.1	14.0	22.3
MTRNet [18]	256	30.96/37.69	90.95/95.83	4.17/1.22			
	512	35.49/40.03	97.10/97.69	2.18/0.80	58.6	13.7	22.2
MTRNet++ [17]	256	37.40/40.26	97.02/97.25	0.86/0.79			
	512	38.31/40.64	97.82/97.94	0.81/0.73	64.0	15.9	25.4
EraseNet [9]	512	34.35/41.73	98.01/98.62	1.81/0.66	30.8	0.6	1.2
EraseNet [9] + M	512	36.37/42.98	98.50/98.75	1.72/0.56	31.4	0.0	1.4
Ours	512	-/43.64	-/98.64	-/0.55	18.9	0.1	0.3

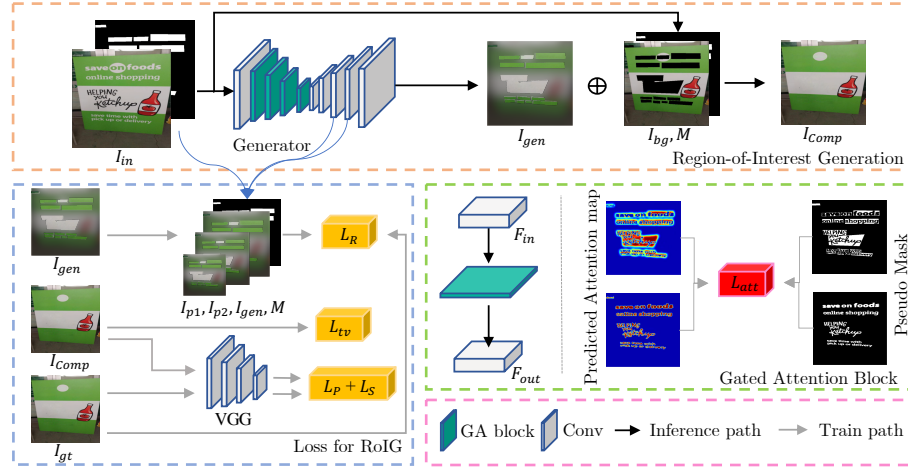


Fig. 2. The overall architecture of the proposed model. The top box, Region-of-Interest Generation, shows that our model generates images focusing only on the text region. The box on the bottom left, Loss for RoIG, how our loss function helps the model focus only on the text region. The L_R only use I_{gen} with M , and L_{tv} , L_p and L_s use I_{comp} , so the outside of text regions in I_{gen} does not participate in loss calculation. On bottom right box, The GA calculates the stroke attention and stroke surrounding attention mask for every layer of the generator’s encoder, then automatically calculates the importance of each. Each attention map is supervised by pseudo mask. More detailed information on GA is in Figure 3.

4 Methodology

4.1 Motivation

Previous STR models [18, 17] used a text box region as well as a text stroke region in an attempt to perform precise text removal. However, TSSR was mostly overlooked. After finding inspiration from observing how humans must alternate between paying attention to the text stroke regions and the surrounding regions of the text while manually performing STR, we devised the GA. Meanwhile, because all previous studies performed STR on the entire image, artifacts frequently occurred in non-text regions. Thus, we devised the RoIG, which allows our STR model to only generate a result image from within the text box region instead of wasting resources attempting to perform STR on the full image.

4.2 Model architecture

Figure 2 shows the architecture of our proposed method. The generator G takes the image and corresponding text box mask as its input and produces a non-text image that is visually plausible. Following GAN-based methods [11, 7], the discriminator D takes both the input of the generator and the target images

as its input and differentiates between real images and images produced by the generator. The objective functions of the generator and discriminator are as follows:

$$L_{adv} = \mathbb{E}_x [\log D(x, G(x))] \quad (1)$$

$$L_D = \mathbb{E}_{x,y} [\log D(x, y)] + \mathbb{E}_x [(1 - \log D(x, G(x)))] \quad (2)$$

where x is the input image concatenated with a text box mask and y is the target image.

Generator. The generator has an FCN-ResNet18 backbone and skip connections [13] between the encoder and decoder. The model is composed of five convolution layers paired with five deconvolution layers with a kernel size of 4, stride of 2, and padding of 1. The convolution pathway is composed of two residual blocks [6], which contains the proposed Gated Attention (GA) module.

Discriminator. For training, we use a local-aware discriminator proposed in EnsNet [22], which only penalizes the erased text patches. The discriminator is trained with locality-dependent labels, which indicate text stroke regions in the output tensor. It guides the discriminator to focus on text regions.

4.3 Gated Attention

Localizing both the TSR and TSSR is imperative to perform surgical STR. In order to do this without letting the size of our model blow up, we used spatial attention[20] instead of a separate image segmentation branch[17]. Table 2 shows how our proposed model is significantly faster and smaller than MTRNet++ [17].

Figure 3 shows the architecture of the GA. The module takes the feature map as its input and generates a TSR and TSSR feature map, then adjusts the proportion of these two feature maps through gate parameters. The process of GA is as follows:

$$F'_i = (\text{MaxPool}(F_i^{In}) \oplus \text{AvgPool}(F_i^{In}) \oplus M_{box}) \quad (3)$$

$$F_i^t = W_i^t \cdot F'_i \quad (4)$$

$$F_i^s = W_i^s \cdot F'_i$$

$$A_i^{out} = \sigma(\alpha_i F_i^t + \beta_i F_i^s) \quad (5)$$

$$F_i^{out} = F_i^{In} A_i^{out} \quad (6)$$

where i denotes the i th layer in the encoder, F^{In} and F^{out} denote the input and output feature maps, W^t and W^s denote 7x7 convolution filters that extract TSR and TSSR features, F^t and F^s denote extracted feature maps for localized TSR and TSSR, α and β denote gate parameters, and σ and A^{out} denote the Sigmoid activation function and the attention score map.

Figure 1 shows the text box mask, TSR mask, and the TSSR mask necessary for this. We generated a pseudo text stroke mask, automatically calculated by taking the pixel value difference between the input image and ground truth image. The TSR masks help train the TSR attention module to distinguish the

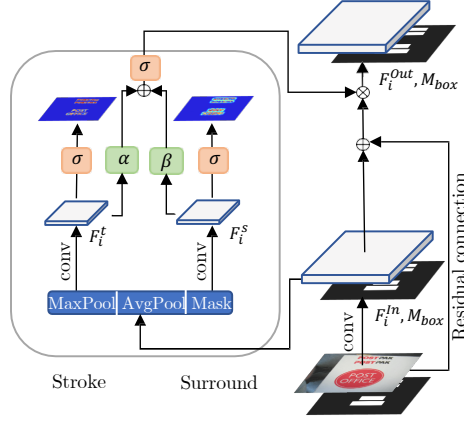


Fig. 3. Architecture of the Gated Attention. The GA calculates the TSR and TSSR attention masks, then automatically calculates the importance of each through gate parameters α, β .

TSR. The TSSR mask is the intersection region between the exterior of the pseudo text stroke mask and the interior of the text box mask. It makes the TSSR attention module train to focus on the colors and textures of the TSSR. Note that the attention method that we used differs from S. Woo *et al.* [20], which was trained in a weakly-supervised manner. We show that our method is superior in Section 5.4. The GA learns the gate parameter on its own, and can thus adjust the respective attention ratios allocated to TSR and TSSR during training. The loss function is designed to be applied only within the text box regions, not to the entire score map. The loss function to train GA is as follows:

$$L_{att}^t = \begin{cases} -Gt_i^t \log(S_i^t) & \text{if } M_i^{Box} > 0, \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

$$L_{att}^s = \begin{cases} -Gt_i^s \log(S_i^s) & \text{if } M_i^{Box} > 0, \\ 0 & \text{otherwise} \end{cases}$$

$$L_{att} = L_{att}^t + L_{att}^s \quad (8)$$

where Gt_i^t and Gt_i^s are the i th pixel of the ground truth mask representing TSR and TSSR. S_i^t and S_i^s are the i th pixel of the TSR and TSSR attention score maps. M_i^{Box} is text box mask. The shape of M , Gt and S is $(H_n, W_n, 1)$. H_n and W_n denote the sizes of feature map in the n th layer.

4.4 Region-of-Interest Generation

Most STR methods attempt to perform in-painting of TSR as well as reconstruction of the entire image. However, our approach can skip the reconstruction of

non-masked regions altogether because if the text box region is given as an input to the STR model, there is no need to render the entire image for the output. Therefore, we modified the loss function so that our model’s generator only has to focus on the text box region during training. Note that the generator’s loss is only calculated with respect to the text box region. Every other region is considered *don’t care* and therefore is irrelevant during training. Because all regions other than the text box obtained from the generator’s output value are not used, we blurred them in Fig 3. This makes training significantly easier.

In total, we modified four loss functions for RoIG: RoI Regression, Perceptual, Style, and Total Variation Loss.

RoI Regression Loss. We modified regression loss to only consider text regions. The proposed *Region-of-Interest Regression Loss* is defined as:

$$L_R(M, I_{out}, I_{gt}) = \sum_{i=1}^n \lambda_i \odot M_i \| I_{out(i)} - I_{gt(i)} \|_1 \quad (9)$$

where $I_{out(i)}$ is the output of the i th deconvolution pathway. We use the output of the 3rd, 4th, and 5th layers in the deconvolution pathway. M_i and $I_{gt(i)}$ are the box mask and ground truth that was resized to the same scale as $I_{out(i)}$. λ_i is the weight for scale. We set λ_i to 0.6, 0.8, 1.0.

Perceptual Loss. Perceptual Loss [8] is used to make the generated image more realistic. It reduces the difference between the high-level features of the two images extracted using the pre-trained ImageNet. Some works [10, 22, 9] used a composited output to address the discrepancy between text-erased regions and the background. We use only the text regions in the generated image by pasting it into the input image. We address the discrepancy between the text box and input images by designing the loss function to use two composited outputs generated using the box and stroke mask, respectively. Perceptual Loss is defined as:

$$\begin{aligned} I_{boxComp} &= I_{in}(1 - M_{Box}) + I_{out}M_{Box} \\ I_{StrokeComp} &= I_{in}(1 - M_{Stroke}) + I_{out}M_{Stroke} \end{aligned} \quad (10)$$

$$\begin{aligned} L_P &= \sum_{n=1}^{N-1} \| A_n(I_{boxComp}) - A_n(I_{gt}) \|_1 \\ &+ \sum_{n=1}^{N-1} \| A_n(I_{StrokeComp}) - A_n(I_{gt}) \|_1 \end{aligned} \quad (11)$$

where I_{in} , I_{out} refer to the input image and the generator’s output image, $I_{BoxComp}$, $I_{StrokeComp}$ are generated images composited of box mask M_{box} and stroke mask M_{stroke} respectively. I_{gt} is the ground truth image, and A_n refers to the activation of the n th layer in network. We use the pool1, pool2 and pool3 layers of the VGG-16 [14] pretrained on ImageNet [3].

Style Loss. Style Loss [4] considers the global texture of the entire image and is used to further improve the visual quality of the output. It is calculated using

the Gram matrix of feature maps. Like perceptual loss, we use two composited outputs. Style loss is defined as:

$$L_S = \sum_{n=1}^{N-1} \left\| \frac{1}{H_n W_n C_n} [\phi(A_n(I_{boxComp})) - \phi(A_n(gt))] \right\|_1 + \sum_{n=1}^{N-1} \left\| \frac{1}{H_n W_n C_n} [\phi(A_n(I_{strokeComp})) - \phi(A_n(gt))] \right\|_1 \quad (12)$$

where $\phi(x) = x^T x$ is a gram matrix operator and A_n is the activation of the n th layers of VGG-16 [14]. We use same layers as Perceptual Loss.

Total Variation Loss. J. Johnson *et al.* [8] proposed total variation loss for global denoising. As our model generates images using RoIG method, the loss function uses composited images which are generated using box masks. The total variation loss is as follows:

$$L_t = \sum_{i,j} \left\| I_{Comp}^{i,j+1} - I_{Comp}^{i,j} \right\|_1 + \left\| I_{Comp}^{i+1,j} - I_{Comp}^{i,j} \right\|_1 \quad (13)$$

where i, j are the pixel positions.

Total Loss function. The total loss function for train generator is as follows:

$$L_G = \lambda_r L_R + \lambda_p L_P + \lambda_s L_S + \lambda_t L_t + \lambda_{adv} L_{adv} + \lambda_{att} L_{Att} \quad (14)$$

where λ_r to λ_{att} represent the weight of each loss. We set λ_i to 100, 0.5, 50.0, 25.0, 1, 10.

5 Experiment

5.1 Dataset and Evaluation Metrics

For training and evaluation, we use both synthetic and real datasets.

Synthetic data. The Oxford Synthetic text dataset [5] is adopted for training and evaluation. The dataset contains around 800,000 images composed of 8,000 text-free images. We randomly selected 95% images for training, selected 10,000 images for testing, and used the rest for validation. Note that the background images in the train set and test set are mutually exclusive.

Real data. SCUT-EnsText [9] is a real dataset for scene text removal. The dataset, which was manually generated from Chinese and English text images, contains 2,749 train and 813 test images. In this paper, we adopted these images for training and evaluation.

Preprocessing. We need stroke-level segmentation masks to train our model. However, the existing datasets do not provide stroke-level segmentation masks. Therefore, we created it automatically by calculating the pixel value difference

between the input image and the ground truth image. To suppress noise, we set a threshold of 25.

We combined synthetic and real datasets. In total, we used 738,113 images for training. The test set was used separately to distinguish between performance on real and synthetic datasets.

Evaluation Metrics T. Nakamura *et al.* [12] proposed an evaluation method using an auxiliary text detector. An auxiliary text detector obtains detection results on the images with text removed. Then, it evaluates the model performance by calculating Precision, Recall, and F-score values. A lower value means that the texts are better erased. In this paper, we use Detection Eval [19] as an evaluation metric and CRAFT [1] as an auxiliary detector. However, that method only indicates how much text has been erased, not output quality. S. Zhang *et al.* [22] proposed using the evaluation method that is used in image inpainting. They used PSNR, SSIM, MSE, AGE, pEPs, pCEPs to evaluate image quality. The higher the value of PSNR and SSIM, and the lower the value of other metrics, the better the quality of the output image. We use PSNR, SSIM, and AGE for evaluation.

5.2 Implementation details

We trained our model for 6 epochs with batch size 30 on the combined dataset. The Adam optimizer with β (0.9, 0.999) was used. We set the initial learning rate to 0.0005 and divided it by 5 every 50,000 steps. PyTorch and NVIDIA Tesla M40 GPUs were used in all experiments.

5.3 Re-implementation details

In this section, we provide details of the re-implementation of previous methods. We re-implemented EnsNet [22] by modifying the code implemented in mxnet and trained the model with the same hyperparameters used in their paper. We re-implemented MTRNet [18] by converting the code implemented in TensorFlow to PyTorch and trained the model with the same batch size and epoch as mentioned in their paper. We used the official implementation of MTRNet++ [17] and EraseNet [9]. We trained them with the same batch size and epoch as mentioned in their respective papers. EraserNet [9] + M refers the model using mask as its input with image.

5.4 Ablation Study

In this section, we validate the effectiveness of our contributions: Gated Attention (GA) and Region-of-Interest Generation (RoIG)

BaseLine. We combined the text region mask with the EnsNet [22] model to make a baseline. This results in a quality improvement of the output image as well as the added functionality of flexibly removing only specific characters at the user’s discretion. All of our experiments including the baseline use a 4-channel input by concatenating the 3-channel RGB and a 1-channel mask.

Table 5. Result quality comparison with ablation studies. SA, TSRA, TSSRA, GA, and RoIG refer to Simple Attention [20], Text Stroke Region Attention, Text Stroke Surrounding Region Attention, Gated Attention, and Region of Interest Generation respectively. The notation is the same as Table 2.

Method	Image Eval			Detection Eval		
	PSNR	SSIM	AGE	P	R	F
Original Image				79.8	67.2	73.0
EnsNet [22]	31.05/32.99	94.78/95.16	2.67/1.85	73.1	54.7	62.6
BaseLine (EnsNet + M)	35.65/37.28	96.67/96.78	1.50/0.89	63.3	33.2	43.6
BaseLine + SA	35.73/37.24	96.53/96.62	1.52/0.89	65.5	34.4	45.1
BaseLine + TSRA	36.07/38.06	97.23/97.34	1.46 /0.84	54.3	21.1	30.4
BaseLine + TSSRA	36.12/38.45	97.29/97.46	1.51/0.84	39.1	6.0	10.5
BaseLine + GA	36.38/38.82	97.56/97.72	1.46/0.80	30.6	3.9	6.9
BaseLine + RoIG	-/40.82	-/98.19	-/0.66	27.5	3.1	5.6
BaseLine + GA + RoIG	-/41.37	-/98.46	-/0.64	15.5	1.0	1.8




Fig. 4. Comparison of the output results after ablation. Image from left to right: Input, Baseline, Baseline + TSRA, Baseline + TSSRA, Baseline + GA, and Baseline + GA + RoIG.

Attention. First, we performed the following three experiments to observe the effect of each attention on the STR results: Simple Attention (SA) [20], Text Stroke Region Attention (TSRA), and Text Stroke Surrounding Region Attention (TSSRA).

Table 5 shows how the application of only SA does not improve the quality of STR. Figure 5’s (a) shows how the application of only SA in STR does not lead to localization of the text stroke and the surrounding regions of the text stroke properly. However, both the TSRA and the TSSRA obtained higher PSNR and SSIM results. In particular, the evaluation result of the TSSRA was significantly better than the TSRA in Detection Eval. This shows that TSSR is more important. Figure 4 demonstrates that results produced by the TSSRA reduce more artifacts than the TSRA. TSRA helps locate the TSR because it is a target for inpainting, but is inappropriate when focusing on the TSSR to erase text. On the other hand, TSSRA is appropriate to focus on the TSSR, and the model utilizes this to fill the TSR and generate higher quality output.

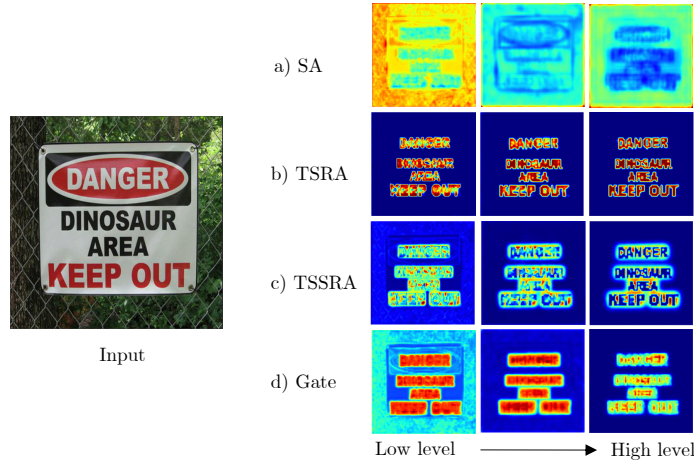


Fig. 5. Visualization of the attention masks for differing encoding layers. The images progressively go from low-level features to high-level features. (a) is a visualization of simple spatial attention [20], (b) and (c) is a visualization of the text stroke and the surrounding attention that the GA generated. (d) is a visualization of how the GA used the gate parameter to aggregate TSA and TSSA. Spatial attention does a poor job finding the text strokes and the surrounding regions if simply applied in STR. In comparison, we can see that the GA pays more attention to the text stroke regions in low-level features while paying more attention to the surrounding regions of the text in high-level features.

Furthermore, We found that having the GA module picking optimal ratios from an ensemble of both TSRA and TSSRA, then aggregating the features as appropriate was super effective. As Table 5 and Figure 4 clearly demonstrate, the evaluation scores of STR were highest when using the GA while also leaving almost no artifacts behind on the resulting image. In Figure 5, the GA module puts more emphasis on the surrounding regions of text strokes rather than the text strokes as it approaches higher-level features.

Region of Interest Generation. In order to measure the effects of RoIG, we performed experiments with and without its use. Table 5 demonstrates that RoIG significantly improves the quality of STR in all metrics. Figure 4 shows that models with the application of RoIG left almost no artifacts with the best overall results.

5.5 Comparison with previous methods

As we mentioned in Section 3, the proposed model maintains a competitive edge with speed and model size while significantly improving the result quality than previous STR models on both real and synthetic data. The first row of Figure 6 shows that EnsNet[22] and EraseNet[9], both of which do not use an explicit text box region, only partially erase text. The fourth row of the Figure shows that




Fig. 6. Comparison of output image results with other prominent STR models. Image from left to right: Input, Ground Truth, EnsNet [22], MTRNet [18], MTRNet++ [17], EraseNet [9], and Ours.

MTRNet[18] and MTRNet++[17] do not successfully remove all text from complex backgrounds without leaving behind artifacts or partially-erased text. However, our proposed model with GA and RoIG successfully outputs high-quality STR images without residual artifacts from images with small text, curved text, and text on complex backgrounds, even without additional refinement.

6 Conclusion

Although there was a lot of progress in the STR area, it was difficult to establish the superiority of a model from the previously proposed methods because there was no standardized and fair way to evaluate performance. In this paper, we re-implemented prominent previously proposed methods, trained and evaluated on respective standardized datasets, and evaluated their accuracies, model size, and inference time in an objectively fair manner. We proposed a simple yet highly impactful STR method with Gated Attention (GA) and Region-of-Interest Generation (RoIG). GA uses attention on the text strokes and the surrounding region’s colors and textures to surgically erase text from images. RoIG makes the generator focus on only the region with text instead of the entire image for more efficient training. Our method significantly outperforms all existing state-of-the-art methods on all benchmark datasets in terms of inference time and output image quality.

Acknowledgements. We wish to thank Osman Tursun for providing codes of MTRNet and MTRNet++.

References

1. Baek, Y., Lee, B., Han, D., Yun, S., Lee, H.: Character region awareness for text detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9365–9374 (2019)
2. Bertalmio, M., Bertozzi, A.L., Sapiro, G.: Navier-stokes, fluid dynamics, and image and video inpainting. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001. vol. 1, pp. I–I. IEEE (2001)
3. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
4. Gatys, L.A., Ecker, A.S., Bethge, M.: A neural algorithm of artistic style. arXiv preprint arXiv:1508.06576 (2015)
5. Gupta, A., Vedaldi, A., Zisserman, A.: Synthetic data for text localisation in natural images. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2315–2324 (2016)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
7. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
8. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European conference on computer vision. pp. 694–711. Springer (2016)
9. Liu, C., Liu, Y., Jin, L., Zhang, S., Luo, C., Wang, Y.: Erasenet: End-to-end text removal in the wild. *IEEE Transactions on Image Processing* **29**, 8760–8775 (2020)
10. Liu, G., Reda, F.A., Shih, K.J., Wang, T.C., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 85–100 (2018)
11. Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
12. Nakamura, T., Zhu, A., Yanai, K., Uchida, S.: Scene text eraser. In: 2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR). vol. 1, pp. 832–837. IEEE (2017)
13. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
14. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
15. Tang, Z., Miyazaki, T., Sugaya, Y., Omachi, S.: Stroke-based scene text erasing using synthetic data for training. *IEEE Transactions on Image Processing* **30**, 9306–9320 (2021)
16. Telea, A.: An image inpainting technique based on the fast marching method. *Journal of graphics tools* **9**(1), 23–34 (2004)
17. Tursun, O., Denman, S., Zeng, R., Sivapalan, S., Sridharan, S., Fookes, C.: Mtr-net++: One-stage mask-based scene text eraser. *Computer Vision and Image Understanding* **201**, 103066 (2020)

18. Tursun, O., Zeng, R., Denman, S., Sivapalan, S., Sridharan, S., Fookes, C.: Mtr-net: A generic scene text eraser. In: 2019 International Conference on Document Analysis and Recognition (ICDAR). pp. 39–44. IEEE (2019)
19. Wolf, C., Jolion, J.M.: Object count/area graphs for the evaluation of object detection and segmentation algorithms. *International Journal on Document Analysis and Recognition* **8**(4), 280–296 (2006)
20. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: Proceedings of the European conference on computer vision (ECCV). pp. 3–19 (2018)
21. Zdenek, J., Nakayama, H.: Erasing scene text with weak supervision. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2238–2246 (2020)
22. Zhang, S., Liu, Y., Jin, L., Huang, Y., Lai, S.: Ensnet: Ensconce text in the wild. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33, pp. 801–808 (2019)