

# Self-Prompting Diffusion Transformer for Open-Vocabulary Scene Text Editing via In-Context Learning

Anonymous Authors<sup>1</sup>

## Abstract

Scene text editing aims to modify text in a target region of an image while preserving surrounding background style and texture. Existing methods rely solely on image background information while neglecting the visual details of target regions, which discards stylistic features in the original text and essentially degrades the task to text rendering. Moreover, the conditions imposed by pre-trained glyph encoders limit the scope of editable text. To address these issues, this paper proposes a self-prompting scene text editing method that constructs style and glyph prompts directly from the original image, without introducing additional style or glyph encoders. We employ a two-stage training strategy: the diffusion transformer is first trained on large-scale self-supervised data and then refined using a small set of paired images. By leveraging the in-context learning capability of the Multi-Modal Diffusion Transformer (MM-DiT), it achieves open-vocabulary and style-consistent text editing. Experimental results on various languages demonstrate that our method achieves the state-of-the-art performance in both text accuracy and style consistency. Our anonymous project page: [mstedit.github.io](https://mstedit.github.io).

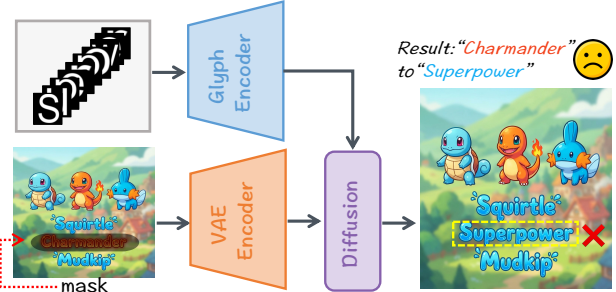
## 1. Introduction

Scene text editing is a specialized yet fundamental image editing task that aims to modify textual content in natural scene images while strictly preserving glyph correctness and visual consistency with the surrounding context. By formulating text modification as targeted content generation within localized text regions, scene text editing naturally aligns with image inpainting, while imposing additional constraints on semantic fidelity, typographic structure, and

<sup>1</sup>Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

(a) Previous OCR-based Text Editing



(b) Self-prompting Text Editing




Figure 1. Comparison of previous OCR-based text edit and our proposed self-prompting text edit.

cross-modal consistency. Driven by recent advances in diffusion-based image inpainting, recent methods (Tuo et al., 2023; 2024; Zeng et al., 2024; Wang et al., 2025a;b) have made notable progress in scene text editing. However, these methods inherit the intrinsic limitations of inpainting-based formulations—they struggle to capture fine-grained glyph structures—particularly for logographic scripts and rare languages—and inevitably discard visual information within the original text regions. As a result, edited outputs often suffer from character distortion and degraded consistency in style and texture with respect to the source text.

The main challenge in scene text editing lies in capturing the glyph features of target text. An intuitive approach is to utilize a pre-trained optical character recognition (OCR) (Du et al., 2020) feature extractor as the glyph encoder. However, as OCR systems are inherently designed as classifiers, their fixed vocabularies fundamentally limit the scope of scene text editing, whereas a glyph encoder trained from scratch relies on massive image data. In addition, existing methods

take the design paradigm of image inpainting for text editing, which discards the original information of the target region. As a consequence, the generated text often fails to preserve the pre-editing style and instead borrows stylistic cues from surrounding regions. As illustrated in Figure 1(a), although previous methods have modified the red word "Charmander" to the target word "Superpower", the target text retains the same blue font as the surrounding text, thereby losing its original stylistic texture.

To address these issues, this paper proposes a self-prompting open-vocabulary text editing method. We capture glyph features at the stroke level rather than the character level, without the need for additional encoders. We leverage the in-context learning capability of the Multi-Modal Diffusion Transformer (MMDiT) (Labs, 2024) to enhance the model's generalization performance. This design enables the model to adapt to complex real-world scenarios and support diverse languages with only a small-scale paired dataset, while simultaneously ensuring strict style consistency between the original and edited text. As demonstrated in Figure 1(b), our method successfully modifies the red word "Charmander" to the red word "Superpower". Notably, the edited text retains the original color, font, and texture of the target region, even though the surrounding text in the background remains blue.

Specifically, for input prompt construction, we generate a high-fidelity glyph map by rendering the target text, which serves as the glyph prompt to guide text content generation. Concurrently, we extract the original pixel information from the target text region in the input image to form a style prompt, which encodes the unique color, texture, and font style of the original text. These two prompts are then concatenated with the full input image to form a multi-modal input tensor for the MMDiT backbone. Second, we freeze the model's encoder and decoder components and exclusively train the backbone using large-scale self-supervised image-text datasets. This step equips the model with fundamental text inpainting capabilities without overfitting to specific text styles. Finally, we utilize a mask-free image editing tool to collect and filter high-quality paired image datasets, where each pair consists of an original image and its corresponding style-consistent edited version. These curated datasets are then used for the cooldown training, during which the model learns to align the generated text with the original style of the target region.

The main contributions of our proposed self-prompting scene text edit can be summarized as follow:

- We achieve open-vocabulary text editing by in-context learning to capture stroke-level features from rendered glyph images, rather than character-level features.
- We enhance pre- and post-editing style consistency

of the target text via cooldown training on a limited amount of paired data, without introducing additional style encoders.

- Extensive experiments on the AnyWord-3M and MST-Edit datasets demonstrate that our method outperforms existing methods in text accuracy and style fidelity across 13 evaluated languages.

## 2. Related Work

### 2.1. Image Inpainting

Image Inpainting aims to fill missing image regions with visually coherent and semantically plausible content while remaining consistent with surrounding context. Early methods rely on patch-based propagation and GAN-based encoder-decoder architectures with structural priors such as partial or gated convolutions and contextual attention (Liu et al., 2018; Yu et al., 2018; 2019). More recent advances are dominated by diffusion-based approaches, which can be categorized into sampling-based methods that modify the denoising process in a training-free manner (Avrahami et al., 2022; Lugmayr et al., 2022), and fine-tuning-based methods that explicitly encode masks and masked images to improve content and shape awareness (Rombach et al., 2022; Manukyan et al., 2023; Zhuang et al., 2024). To balance performance and generality, recent plug-and-play designs decouple masked image conditioning from generation, enabling effective inpainting without retraining entire diffusion backbones (Zhang et al., 2023; Ju et al., 2024).

Image inpainting serves as the foundation of scene text editing, where accurate background restoration is crucial for consistent text synthesis.

### 2.2. Scene Text Editing

Scene text editing extends image inpainting by jointly requiring background restoration and geometry- and style-consistent text synthesis. Early methods (Roy et al., 2020; Yang et al., 2020; Qu et al., 2023) adopt multi-stage GAN-based pipelines with explicit geometric or stroke priors, offering controllable edits but suffering from poor generalization and error accumulation in complex scenes. Recent methods (Tuo et al., 2023; 2024; Wang et al., 2025b) employ pre-trained OCR models as glyph encoders to capture character-level structures. While this design improves glyph awareness, it inherently restricts generalization to characters outside the fixed OCR vocabulary. In contrast, Flux-Text (Lan et al., 2025) and TextFlux (Xie et al., 2025) remove glyph encoders and instead render glyphs directly within masked regions to extract visual features; however, the spatial extent of the target region can hinder the accurate capture fine-grained glyph details. Despite their differences, these methods largely adhere to an image inpainting formu-

lation that discards the original text region, leading to the loss of pre-existing style information and degraded stylistic consistency during editing. TextCtrl (Zeng et al., 2024) employs a text style encoder to capture color, font, texture, and background features from the original text. However, the representation mismatch between the glyph-structure encoder and the VAE image latent space restricts the method’s scalability beyond isolated text regions.

Differing from existing methods, our method directly constructs style and glyph prompts from the original image without additional encoders, thereby enabling open-vocabulary and style-consistent text editing.

### 3. Method

#### 3.1. Preliminary

Our method is built upon *FLUX-Fill* (Labs, 2024), an inpainting-oriented variant of MMDiT, which formulates image editing as a conditional rectified flow process in the latent space and employs a transformer-based architecture to jointly reason over visual content, spatial structure, and textual semantics under multimodal conditions.

**Masked Image Construction.** Given an input image  $I \in \mathbb{R}^{H \times W \times 3}$  and a binary mask  $M \in \{0, 1\}^{H \times W}$  indicating the target text region, FLUX-Fill first constructs a masked image

$$I_m = I \odot (1 - M), \quad (1)$$

where  $\odot$  denotes element-wise multiplication. The masked image  $I_m$  preserves the visual context outside the editable region while removing the original content within the target area. Both the original image  $I$  and the masked image  $I_m$  are encoded into latent representations using a frozen VAE encoder, producing visual tokens that serve as the input to the diffusion transformer. The binary mask  $M$  is also embedded and provided to the model as an explicit spatial prior, enabling region-aware generation during denoising.

**Text Prompt Encoding.** FLUX-Fill employs two complementary text encoders to extract semantic guidance from textual prompts. Specifically, a T5 encoder is used to process the full natural-language prompt, capturing high-level semantic intent and contextual information. The resulting embeddings are injected into the diffusion transformer through cross-attention layers, guiding global content generation and scene-level consistency. In parallel, a CLIP text encoder is used to encode concise textual descriptors that emphasize visual alignment, such as object names or short phrases. CLIP embeddings primarily serve to enhance vision–language alignment and stabilize the correspondence between generated content and visual context.

**Dual-Stream and Single-Stream Transformer.** The core of FLUX-Fill is a hybrid transformer design that alternates

between dual-stream and single-stream blocks. In the dual-stream transformer blocks, visual tokens and text tokens are processed in separate streams with independent self-attention operations. Cross-attention is then applied to enable information exchange between modalities, allowing the model to align textual semantics with spatially grounded visual representations while maintaining modality-specific feature structures.

After cross-modal interaction, the architecture transitions to single-stream transformer blocks, where visual and text tokens are concatenated into a unified token sequence and processed jointly.

**Rectified Flow Objective.** Let  $z_0$  denote the clean latent representation obtained from the VAE encoder and  $z_1 \sim \mathcal{N}(0, I)$  be Gaussian noise. A noisy latent  $z_t$  is constructed via rectified flow interpolation at timestep  $t$ , and the model is trained to predict the velocity field connecting  $z_t$  and  $z_0$ :

$$\mathcal{L}_{\text{RF}} = \mathbb{E}_{t, z_0, z_1} \left[ \|\hat{v}_\theta(z_t, t, c) - (z_1 - z_0)\|_2^2 \right], \quad (2)$$

where  $c$  denotes the multimodal conditioning inputs, including visual tokens, text embeddings, and spatial mask information.

#### 3.2. Overall Architecture

Our method explicitly disentangles text style and text content through dedicated visual and textual prompts, while fully leveraging the in-context learning capability of the MMDiT backbone. An overview of the proposed architecture is shown in Figure 2.

**Style Prompt Construction.** To preserve the visual appearance of the original text, we construct a *style prompt* from both visual and textual perspectives.

Given an input image  $I \in \mathbb{R}^{H \times W \times 3}$  and a binary mask  $M \in \{0, 1\}^{H \times W}$  indicating the target text region, we compute the maximal enclosing bounding rectangle of the masked area and crop the corresponding region from  $I$  to obtain the visual style prompt  $I_s$ . This cropped patch encodes region-specific appearance information, such as color, texture, font characteristics, and local illumination, and serves as a visual reference for style preservation.

For style text prompt, we encode the input text description using the CLIP text encoder. Together, the visual and textual style prompts enable the model to infer and preserve the intrinsic style of the target text region without introducing additional style-specific encoders.

**Glyph Prompt Construction.** To guide the generation of target text content, we construct a *glyph prompt* that explicitly represents the desired textual structure.

Specifically, the target text is rendered into a single-line

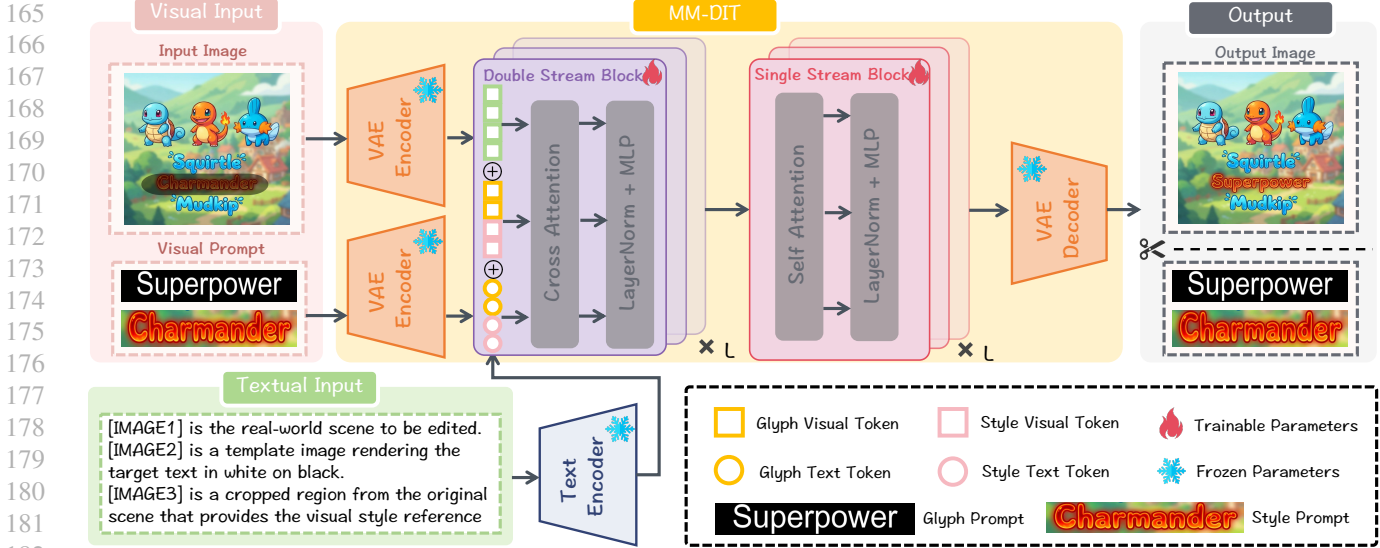


Figure 2. Overview of our proposed method.

glyph image using the Pillow library, producing a white-on-black glyph map  $I_g$ . This high-contrast representation preserves fine-grained stroke-level geometry and provides explicit structural guidance for complex glyphs across different languages and scripts, reducing the burden on the diffusion model to learn character structures from scratch.

For glyph text prompt, we encode the target text string using the T5 text encoder. The resulting embeddings capture high-level semantic and syntactic information of the target text and are injected into the MMDiT backbone through cross-attention, guiding content-aware generation.

**Denoising Process.** The visual glyph prompt  $I_g$ , visual style prompt  $I_s$ , and masked image  $I_m$  are concatenated along the channel dimension to form the composite visual input

$$I_{\text{input}} = \text{Concat}(I_g, I_s, I_m). \quad (3)$$

The composite input is encoded by a frozen VAE encoder  $\mathcal{E}(\cdot)$  to obtain the latent representation  $z_0$ , which is processed by the MMDiT backbone together with the textual glyph and style embeddings.

After the denoising process, the predicted latent representation is decoded by the VAE decoder  $\mathcal{D}(\cdot)$  to produce an edited image. Since the masked image is used only for conditioning, the final output is obtained by cropping the decoded result to the spatial region corresponding to the target text area.

### 3.3. Cooldown Training

Existing OCR-based datasets discard original text appearance by masking target regions, providing only self-supervised rendering signals and limiting style preservation.

To enable style-consistent text editing under limited paired data, we adopt a two-stage cooldown training strategy.

We first construct a paired image dataset tailored for style-aware learning. Specifically, we leverage an instruction-based image editing model (*Nano Banana Pro*<sup>1</sup>) to generate edited images conditioned on explicit editing instructions. Each data pair consists of an original image and a corresponding edited image, in which only the target text region is modified while all other regions are strictly preserved. To ensure data reliability, we manually filter the generated pairs and retain only samples that satisfy the following criteria:

- Non-target regions remain unchanged.
- The generated text content is semantically accurate.
- The edited text preserves consistent style, color, and texture with respect to the original text.

**Self-supervised Pretraining.** We first pretrain the model on the AnyWord-3M dataset, which provides large-scale self-supervised data for multilingual scene text rendering. During this stage, the optimization objective follows Eq. 2, where the conditioning includes multimodal visual and textual inputs.

**Cooldown Training.** Inspired by MECO (Gao et al., 2025), we treat the original text region in the input image as meta-information during training. We continue training from the pretrained checkpoint using a curated paired dataset consisting of 4,000 manually filtered image pairs. Each pair contains an original image and a corresponding edited image with style-consistent text replacement. The comparison

<sup>1</sup><https://nanobanana.im/nano-banana-pro>

between data used by pre-training and cooldown stage is shown in Figure 3

To mitigate the degenerate optimization where original text region features unduly dominate the target region (leading the model to replicate both original style and glyphs), we design a target-region-oriented training objective for the cooldown stage. This objective restricts the learning signal to localized text transformations, prompting the model to decouple style preservation from content regeneration in optimization.

Let  $z_0^{\text{src}}$  and  $z_0^{\text{tgt}}$  denote the latent representations of the source image and its corresponding edited image, respectively. We construct an interpolated latent variable

$$z_t = (1 - \sigma_t) z_0^{\text{src}} + \sigma_t z_0^{\text{tgt}}, \quad (4)$$

where  $\sigma_t \in [0, 1]$  is a time-dependent interpolation coefficient. The model is trained to predict the velocity field that connects the source and target latents by minimizing

$$\mathcal{L}_{\text{CD}} = \mathbb{E}_t \left[ \left\| \hat{v}_\theta(z_t, t, c) - (z_0^{\text{tgt}} - z_0^{\text{src}}) \right\|_2^2 \right], \quad (5)$$

where the conditioning  $c$  includes the original text region as meta-information through the style prompt.

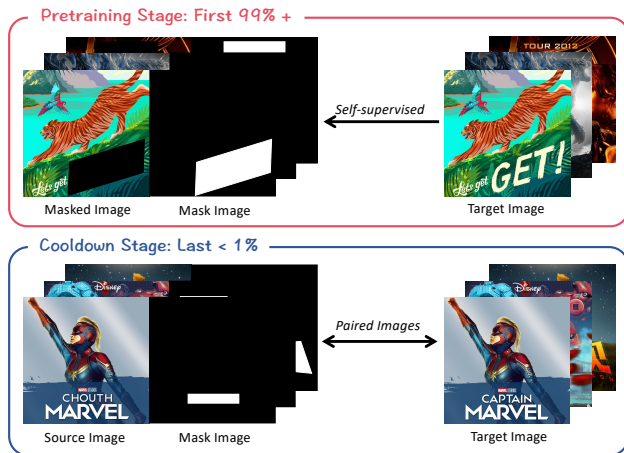


Figure 3. A comparison between data used by standard pre-training and cooldown training.

## 4. Experiment

### 4.1. Experimental Setup

**Datasets.** We adopt AnyWord-3M (Tuo et al., 2023) as the large-scale benchmark dataset. Its training set contains 1.6M Chinese images, 1.39M English images, and 10K images spanning Japanese, Korean, Arabic, Bengali, and Hindi. The test split includes 1,000 Chinese and 1,000 English images, and is denoted as AnyText-benchmark.

Table 1. Composition of the MST-Edit dataset.

Source Dataset	Language	# Images
ICDAR-19	Arabic	1,000
	French	1,000
	German	1,000
	Korean	1,000
	Japanese	1,000
	Italian	1,000
	Bengali	1,000
	Hindi	1,000
RusTitW	Russian	3,795
ThaiOCRBench	Thai	2,808
Swahili-STR	Swahili	985

We further construct a Multi-lingual Scene Text Editing Dataset (MST-Edit) by aggregating multiple publicly available multilingual datasets, including ICDAR-19 (Nayef et al., 2019), RusTitW (Markov et al., 2023), ThaiOCRBench (Nonesung et al., 2025), and Swahili-STR (Douamba et al., 2024). Detailed statistics of the dataset composition are summarized in Table 1. We randomly split 20% of MST-Edit for testing, with the remainder used for training.

**Evaluation metrics.** For textual accuracy, Sentence Accuracy (Seq. ACC) quantifies the correctness of generated text at the sentence level, and Normalized Edit Distance (NED) measures character-level similarity between the generated and target text. For visual fidelity, Fréchet Inception Distance (FID) (Seitzer, 2020) quantifies the distribution alignment of generated and real images in the Inception-v3 feature space, and Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al., 2018) measures the L2 distances between perceptual VGG-based deep features.

**Implementation Details.** All experiments are conducted on a cluster with 8 NVIDIA A100 GPUs. Our model is initialized from FLUX.1-Fill-Dev<sup>2</sup>. During training, we freeze the VAE, CLIP text encoder, and T5 text encoder, and update only the transformer parameters. We follow the default FLUX configuration, using a fixed guidance scale of 30 and 30 sampling steps for both training and inference.

Training is conducted in two stages. We first train the model on AnyWord-3M for one epoch, followed by a 10-epoch cooldown phase on the paired image dataset. Multilingual data are jointly mixed without language-specific scheduling. We use the AdamW optimizer with a constant learning rate of  $2 \times 10^{-5}$ , bf16 mixed precision, and 8-bit optimizer states. The per-GPU batch size is 1 with gradient accumulation over 8 steps, resulting in an effective batch size of 64, and all experiments are conducted with a fixed random seed of 42.

<sup>2</sup><https://huggingface.co/black-forest-labs/FLUX.1-Fill-dev>

Table 2. Quantitative results on AnyText-benchmark. **Bold** indicates the best result and underline indicates the second best.

Method	English				Chinese			
	Sen. Acc(↑)	NED(↑)	FID(↓)	LPIPS(↓)	Sen. Acc(↑)	NED(↑)	FID(↓)	LPIPS(↓)
TextDiffuser (Chen et al., 2023)	0.5176	0.7618	29.76	0.1564	0.0559	0.1218	34.19	0.1252
AnyText (Tuo et al., 2023)	0.6843	0.8588	21.59	0.1106	0.6476	0.8210	20.01	0.0943
TextCtrl (Zeng et al., 2024)	0.5853	0.8146	35.73	0.1978	0.3580	0.6084	49.79	0.2298
AnyText2 (Tuo et al., 2024)	0.7915	0.9100	29.76	0.1734	0.7022	0.8420	26.52	0.1444
GlyphMastero (Wang et al., 2025a)	0.8170	-	-	-	0.7301	-	-	-
FIUX-Text (Lan et al., 2025)	0.8175	0.9193	<u>12.35</u>	<u>0.0674</u>	0.7213	0.8555	<u>12.41</u>	<u>0.0487</u>
TextFlux (Xie et al., 2025)	<u>0.8231</u>	<u>0.9235</u>	13.42	0.0721	<u>0.7289</u>	<u>0.8612</u>	13.67	0.0524
Our method	<b>0.8857</b>	<b>0.9568</b>	<b>7.62</b>	<b>0.0365</b>	<b>0.8249</b>	<b>0.9147</b>	<b>7.95</b>	<b>0.0268</b>

### 4.2. Quantitative Result

We perform quantitative experiments under both bilingual (Chinese/English) and multilingual (non-Chinese/English) settings. For Chinese and English, we evaluate our method against state-of-the-art methods (Chen et al., 2023; Tuo et al., 2023; 2024; Zeng et al., 2024; Wang et al., 2025a; Xie et al., 2025; Lan et al., 2025) on the AnyText-benchmark.

For non-Chinese/English languages, we re-implement two OCR-free baselines (Xie et al., 2025; Lan et al., 2025) and compare them with our method on the MST-Edit dataset, which covers the remaining 11 languages.

Table 2 reports the quantitative results on the AnyText benchmark. Based on these results, we make the following observations:

- Overall performance.** Our method achieves the best performance across all metrics on both the English and Chinese subsets, consistently outperforming existing methods.
- Text accuracy.** Compared with the second-best method, TextFlux, our approach improves sentence-level accuracy from 0.8231 to 0.8857 on English and from 0.7289 to 0.8249 on Chinese. The NED score is also increased from 0.9235 to 0.9568 and from 0.8612 to 0.9147, respectively, indicating improved preservation of fine-grained glyph structures.
- Image quality.** More significant gains are observed on image quality metrics. Our method reduces FID to 7.62/7.95 and LPIPS to 0.0365/0.0268 on English/Chinese, surpassing the strongest baselines and demonstrating superior visual fidelity and stylistic consistency.

The experimental results on MST-Edit are illustrated in Figure 4. Overall, our method consistently outperforms the two OCR-free baselines across all evaluated languages, demonstrating a clear method-level advantage. From a cross-lingual perspective, languages belonging to the Latin script family (e.g., French, German, and Italian) generally achieve better performance, which can be attributed to their shared alphabetic structures and similar stroke patterns. In contrast, languages with more complex glyph compositions and di-




Figure 4. Quantitative comparison of OCR-free methods on MST-Edit, with Seq. ACC on the Y-axis and FID on the X-axis.

verse stroke layouts, such as Thai, tend to exhibit relatively lower performance. These results suggest that character structural complexity plays an important role in multilingual scene text editing, while our method remains robust across diverse writing systems.

### 4.3. Ablation Study

We conduct ablation studies to analyze the effect of style prompts and to investigate potential cross-lingual negative interference in multilingual scene text editing.

To analyze the impact of style prompts, we conduct an ablation study by removing visual style prompts from the image input and textual style prompts from the text input, while disabling the cooldown training described in Section 3.3.

Results on the AnyText benchmark are reported in Table 3. With the cooldown stage enabled, FID is reduced from 11.56 to 7.62 on English and from 12.01 to 7.95 on Chinese, while LPIPS decreases from 0.0608 to 0.0365 and from 0.0425

Table 3. Ablation study on the effect of the cooldown stage.

	Metric	w/o Cooldown	w/ Cooldown
English	Sen. Acc ( $\uparrow$ )	0.8738	<b>0.8857</b>
	NED ( $\uparrow$ )	0.9470	<b>0.9568</b>
	FID ( $\downarrow$ )	11.56	<b>7.62</b>
	LPIPS ( $\downarrow$ )	0.0608	<b>0.0365</b>
Chinese	Sen. Acc ( $\uparrow$ )	0.8125	<b>0.8249</b>
	NED ( $\uparrow$ )	0.8906	<b>0.9147</b>
	FID ( $\downarrow$ )	12.01	<b>7.95</b>
	LPIPS ( $\downarrow$ )	0.0425	<b>0.0268</b>

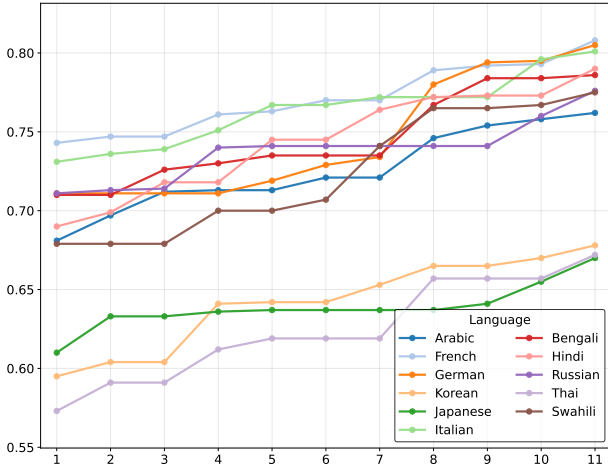


Figure 5. Effect of the number of introduced languages on Seq. ACC of initially introduced languages.

to 0.0268, respectively. In contrast, improvements in text accuracy are relatively modest. These results indicate that style prompts, reinforced by the cooldown stage, primarily improve image quality, while offering complementary gains in text accuracy.

We investigate the impact of progressively incorporating additional language data on earlier-introduced languages using a cyclic training protocol, in which languages are sequentially added in the order *Arabic, English, French, Chinese, German, Korean, Japanese, Italian, Bengali, Hindi, Russian, Thai, Swahili*.

Empirically, Seq. ACC remains stable and exhibits a consistent upward trend as additional languages are introduced, as illustrated in Figure 5. This behavior is consistent with our design principle of learning stroke-level visual primitives rather than language- and character-specific representations: shared low-level stroke structures across writing systems are reinforced by multilingual exposure, leading to improved generalization.

Figure 6 presents qualitative image editing results with and without style prompts. It can be observed that both settings



Figure 6. Comparison of scene text edit results with and without style prompts.

accurately generate the target text without erroneous strokes. However, in the left example, when the style prompt is removed, the target text “HOT SUMMER” is rendered using a font style that does not appear in the original image, leading to noticeable stylistic inconsistency. In the right example, although the result without style prompts renders the target text “意识” using a font consistent with the surrounding text, it fails to preserve the original text background. In contrast, the setting with style prompts produces target text that more faithfully restores both the original font style and background, resulting in higher visual fidelity.

#### 4.4. Qualitative Result

We conduct a qualitative comparison with representative multilingual scene text editing methods, including AnyText2 (Tuo et al., 2024), TextFlux (Xie et al., 2025), and FluxText (Labs, 2024).

As shown in Figure 7, our method consistently generates accurate target text across different scripts, preserving correct character structures and visual consistency. More specifically, results on Russian, Japanese, and Thai demonstrate



Figure 7. Qualitative results across Chinese, English, Korean, Japanese, Thai, and Russian.

the effectiveness of our method on long-tail languages with complex and diverse writing systems. In particular, our approach better preserves the original text style in the “战友定制” example, and remains robust in the “GOOD DOG!” case, where the masked region extends beyond the target text, introducing minimal unintended modifications to non-target regions. These observations indicate that our method generalizes well across languages while enabling precise and controlled text editing under challenging conditions.

## 5. Conclusion

This paper presents a self-prompting diffusion transformer framework for open-vocabulary scene text editing that preserves both textual correctness and visual style consistency

by exploiting in-context information from the original image. By directly constructing glyph and style prompts from the input image, the proposed approach enables coherent text generation that remains faithful to the surrounding font appearance, layout, and visual context, without introducing additional glyph or style encoders. Extensive experiments on large-scale multilingual benchmarks demonstrate that our approach achieves strong performance across diverse languages and scripts, while maintaining robustness under progressive language expansion without negative transfer. Overall, this paper provides a unified and scalable solution for extending style-consistent editing to more complex and diverse real-world scenarios.

## Impact Statement

This work studies scene text editing, which aims to modify textual content in images while preserving visual style and background consistency. The proposed method advances the flexibility and robustness of text editing across languages and scripts, which may benefit applications such as graphic design, content creation, and multilingual visual communication.

At the same time, like other image editing and generative techniques, scene text editing could be misused to alter visual content in misleading or deceptive ways. This work is intended for research purposes, and we encourage responsible use in accordance with existing ethical guidelines for generative models and visual media editing.

## References

- Avrahami, O., Lischinski, D., and Fried, O. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18208–18218, 2022.
- Chen, J., Huang, Y., Lv, T., Cui, L., Chen, Q., and Wei, F. Textdiffuser: Diffusion models as text painters. *Advances in Neural Information Processing Systems*, 36: 9353–9387, 2023.
- Douamba, F. W., Song, J., Fu, L., Liu, Y., and Bai, X. The first swahili language scene text detection and recognition dataset. In *International Conference on Document Analysis and Recognition*, pp. 215–226. Springer, 2024.
- Du, Y., Li, C., Guo, R., Yin, X., Liu, W., Zhou, J., Bai, Y., Yu, Z., Yang, Y., Dang, Q., et al. Pp-ocr: A practical ultra lightweight ocr system. *arXiv preprint arXiv:2009.09941*, 2020.
- Gao, T., Wettig, A., He, L., Dong, Y., Malladi, S., and Chen, D. Metadata conditioning accelerates language model pre-training. *arXiv preprint arXiv:2501.01956*, 2025.
- Ju, X., Liu, X., Wang, X., Bian, Y., Shan, Y., and Xu, Q. Brushnet: A plug-and-play image inpainting model with decomposed dual-branch diffusion. In *European Conference on Computer Vision*, pp. 150–168. Springer, 2024.
- Labs, B. F. Flux. <https://github.com/black-forest-labs/flux>, 2024.
- Lan, R., Bai, Y., Duan, X., Li, M., Jin, D., Xu, R., Sun, L., and Chu, X. Flux-text: A simple and advanced diffusion transformer baseline for scene text editing. *arXiv preprint arXiv:2505.03329*, 2025.
- Liu, G., Reda, F. A., Shih, K. J., Wang, T.-C., Tao, A., and Catanzaro, B. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 85–100, 2018.
- Lugmayr, A., Danelljan, M., Romero, A., Yu, F., Timofte, R., and Van Gool, L. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11461–11471, 2022.
- Manukyan, H., Sargsyan, A., Atanyan, B., Wang, Z., Navasardyan, S., and Shi, H. Hd-painter: high-resolution and prompt-faithful text-guided image inpainting with diffusion models. In *The Thirteenth International Conference on Learning Representations*, 2023.
- Markov, I., Nesteruk, S., Kuznetsov, A., and Dimitrov, D. Rustitw: Russian language text dataset for visual text in-the-wild recognition. *arXiv preprint arXiv:2303.16531*, 2023.
- Nayef, N., Patel, Y., Busta, M., Chowdhury, P. N., Karatzas, D., Khlif, W., Matas, J., Pal, U., Burie, J.-C., Liu, C.-I., et al. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition—rrc-mlt-2019. In *2019 International conference on document analysis and recognition (ICDAR)*, pp. 1582–1587. IEEE, 2019.
- Nonesung, S., Jaknamon, T., Chaiophat, S., Nitarach, N., Wittayasakpan, C., Sirichotedumrong, W., Na-Thalang, A., and Pipatanakul, K. Thaiocrbench: A task-diverse benchmark for vision-language understanding in thai. *arXiv preprint arXiv:2511.04479*, 2025.
- Qu, Y., Tan, Q., Xie, H., Xu, J., Wang, Y., and Zhang, Y. Exploring stroke-level modifications for scene text editing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 2119–2127, 2023.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Roy, P., Bhattacharya, S., Ghosh, S., and Pal, U. Steffann: scene text editor using font adaptive neural network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13228–13237, 2020.
- Seitzer, M. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, August 2020. Version 0.3.0.
- Tuo, Y., Xiang, W., He, J.-Y., Geng, Y., and Xie, X. Anytext: Multilingual visual text generation and editing. *arXiv preprint arXiv:2311.03054*, 2023.

- 495 Tuo, Y., Geng, Y., and Bo, L. Anytext2: Visual text gen-  
496 eration and editing with customizable attributes. *arXiv*  
497 *preprint arXiv:2411.15245*, 2024.
- 498 Wang, T., Liu, T., Qu, X., Wu, C., Liu, L., and Hu, X. Glyph-  
499 mastero: A glyph encoder for high-fidelity scene text edit-  
500 ing. In *Proceedings of the Computer Vision and Pattern*  
501 *Recognition Conference*, pp. 28523–28532, 2025a.
- 503 Wang, Y., Zhang, W., Xu, H., and Jin, C. Dreamtext: High  
504 fidelity scene text synthesis. In *Proceedings of the Com-*  
505 *puter Vision and Pattern Recognition Conference*, pp.  
506 28555–28563, 2025b.
- 508 Xie, Y., Zhang, J., Chen, P., Wang, Z., Wang, W., Gao, L.,  
509 Li, P., Sun, H., Zhang, Q., Qiao, Q., et al. Textflux: An  
510 ocr-free dit model for high-fidelity multilingual scene text  
511 synthesis. *arXiv preprint arXiv:2505.17778*, 2025.
- 513 Yang, Q., Huang, J., and Lin, W. Swaptext: Image based  
514 texts transfer in scenes. In *Proceedings of the IEEE/CVF*  
515 *Conference on Computer Vision and Pattern Recognition*,  
516 pp. 14700–14709, 2020.
- 517 Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., and Huang, T. S.  
518 Generative image inpainting with contextual attention. In  
519 *Proceedings of the IEEE conference on computer vision*  
520 *and pattern recognition*, pp. 5505–5514, 2018.
- 522 Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., and Huang, T. S.  
523 Free-form image inpainting with gated convolution. In  
524 *Proceedings of the IEEE/CVF international conference*  
525 *on computer vision*, pp. 4471–4480, 2019.
- 527 Zeng, W., Shu, Y., Li, Z., Yang, D., and Zhou, Y. Textctrl:  
528 Diffusion-based scene text editing with prior guidance  
529 control. *Advances in Neural Information Processing Sys-*  
530 *tems*, 37:138569–138594, 2024.
- 532 Zhang, L., Rao, A., and Agrawala, M. Adding conditional  
533 control to text-to-image diffusion models. In *Proceedings*  
534 *of the IEEE/CVF international conference on computer*  
535 *vision*, pp. 3836–3847, 2023.
- 537 Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang,  
538 O. The unreasonable effectiveness of deep features as a  
539 perceptual metric. In *Proceedings of the IEEE conference*  
540 *on computer vision and pattern recognition*, pp. 586–595,  
541 2018.
- 542 Zhuang, J., Zeng, Y., Liu, W., Yuan, C., and Chen, K. A  
543 task is worth one word: Learning with task prompts for  
544 high-quality versatile image inpainting. In *European*  
545 *Conference on Computer Vision*, pp. 195–211. Springer,  
546 2024.
- 547  
548  
549

550 **A. You *can* have an appendix here.**

551 You can have as much text here as you want. The main body must be at most 8 pages long. For the final version, one more  
552 page can be added. If you want, you can use an appendix like this one.  
553

554 The `\onecolumn` command above can be kept in place if you prefer a one-column appendix, or can be removed if you  
555 prefer a two-column appendix. Apart from this possible change, the style (font size, spacing, margins, page numbering, etc.)  
556 should be kept the same as the main body.  
557

558  
559  
560  
561  
562  
563  
564  
565  
566  
567  
568  
569  
570  
571  
572  
573  
574  
575  
576  
577  
578  
579  
580  
581  
582  
583  
584  
585  
586  
587  
588  
589  
590  
591  
592  
593  
594  
595  
596  
597  
598  
599  
600  
601  
602  
603  
604