

MDiff4STR: Mask Diffusion Model for Scene Text Recognition

Yongkun Du^{1*}, Miaomiao Zhao^{2*}, Songlin Fan^{1,3}, Zhineng Chen^{1†}, Caiyan Jia², Yu-Gang Jiang¹

¹Institute of Trustworthy Embodied AI, Fudan University, China

²School of Computer Science and Technology, Beijing Jiaotong University, China

³China Mobile Shanghai ICT Co., Ltd., China

ykdu23@m.fudan.edu.cn, {zhinchen,slfan,ygj}@fudan.edu.cn, {miaomiao_zhao,cyjia}@bjtu.edu.cn

Abstract

Mask Diffusion Models (MDMs) have recently emerged as a promising alternative to auto-regressive models (ARMs) for vision-language tasks, owing to their flexible balance of efficiency and accuracy. In this paper, for the first time, we introduce MDMs into the Scene Text Recognition (STR) task. We show that vanilla MDM lags behind ARMs in terms of accuracy, although it improves recognition efficiency. To bridge this gap, we propose MDiff4STR, a Mask Diffusion model enhanced with two key improvement strategies tailored for STR. Specifically, we identify two key challenges in applying MDMs to STR: noising gap between training and inference, and overconfident predictions during inference. Both significantly hinder the performance of MDMs. To mitigate the first issue, we develop six noising strategies that better align training with inference behavior. For the second, we propose a token-replacement noise mechanism that provides a non-mask noise type, encouraging the model to reconsider and revise overly confident but incorrect predictions. We conduct extensive evaluations of MDiff4STR on both standard and challenging STR benchmarks, covering diverse scenarios including irregular, artistic, occluded, and Chinese text, as well as whether the use of pretraining. Across these settings, MDiff4STR consistently outperforms popular STR models, surpassing state-of-the-art ARMs in accuracy, while maintaining fast inference with only three denoising steps.

Code — <https://github.com/Topdu/OpenOCR>

1 Introduction

Scene Text Recognition (STR), as a base task in Optical Character Recognition (OCR) systems, has remained a focal point of research in computer vision. In natural scenes, STR faces a wide range of complex challenges (Chen et al. 2022), including curved and distorted text, varying orientations, occlusions, image blur, and artistic fonts. To address these issues, researchers have proposed a wealth of innovative solutions that significantly enhance the robustness and accuracy of recognition models in real-world applications.

Among STR methods, auto-regressive models (ARMs) (Shi et al. 2019; Sheng, Chen, and Xu 2019; Li et al. 2019;

*These authors contributed equally.

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

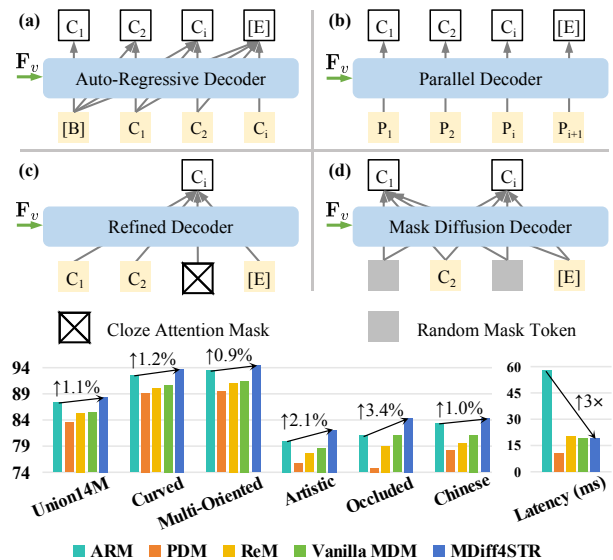


Figure 1: (a) Auto-regressive models (ARMs), (b) Parallel decoding models (PDMs), (c) BERT-like refinement models (ReMs), (d) Mask diffusion models (MDMs). F_v means visual features. MDMs learn to reconstruct character sequences from partially masked inputs through a denoising process, capturing more flexible and comprehensive omnidirectional dependencies than ARMs and refinement models.

Yue et al. 2020; Jiang et al. 2023; Xie et al. 2022; Zheng et al. 2023, 2024; Xu et al. 2024; Yang et al. 2024; Zhou et al. 2024; Du et al. 2025c,a) have emerged as one of the most prominent due to their strong sequence modeling capabilities and achieved state-of-the-art results across standard and challenging benchmarks (Wang et al. 2021; Jiang et al. 2023; Chen et al. 2021). However, the inherently sequential nature of ARMs limits their decoding efficiency.

Recently, mask diffusion models (MDMs) (Shi et al. 2024; Sahoo et al. 2024; Nie et al. 2025; You et al. 2025), a novel non-auto-regressive paradigm, have demonstrated superior performance in both efficiency and accuracy. On the one hand, as shown in Fig. 1(d), MDMs learn to reconstruct original sequences from partially masked inputs by progressively denoising them. This formulation not only over-

comes the unidirectional (left-to-right) modeling limitation of ARMs, but also surpasses the rigid bidirectional modeling of BERT-like models, enabling the capture of more flexible and comprehensive omnidirectional dependencies (Shi et al. 2024; Sahoo et al. 2024). Given that STR is fundamentally reliant on strong language understanding, MDMs show great potential in offering a novel and promising paradigm for enhancing STR performance. On the other hand, its denoising process is highly efficient and controllable, allowing accurate predictions to be produced within only a few steps.

In this paper, we introduce, for the first time, MDMs into the STR task. However, experimental results reveal that the vanilla MDM offers notable advantages in inference efficiency, but it still falls short in recognition accuracy compared to ARMs. To uncover the causes of this performance gap, we identify two key limitations: (1) Noising gap between training and inference. MDMs are typically trained with randomly noisy, while during inference, the model is exposed to structured and deterministic noisy patterns. These patterns are rarely encountered during training, resulting in poor generalization and degraded recognition performance; (2) Overconfident predictions during inference. We observe that MDMs tend to assign excessively high confidence scores to their predictions even when incorrect. This overconfidence hampers the effectiveness of the confidence-based remask mechanism, making it difficult for the model to identify and correct earlier mistakes, thereby undermining the overall efficacy of the multi-step denoising process.

To address the aforementioned challenges, we propose MDiff4STR, a Mask Diffusion Model tailored for STR. It is a novel MDM paradigm for STR incorporating two key innovations. First, we introduce six noising strategies for training. They accurately simulate the noising patterns encountered during inference, effectively mitigating the noising gap. Second, we introduce a token-replacement noise mechanism, a novel noise type distinct from masking. This mechanism enables the model to reconsider and correct its own overconfident yet incorrect predictions, leading to more accurate recognition results.

Experiments on multiple public STR benchmarks demonstrate that MDiff4STR consistently outperforms popular STR models, surpassing state-of-the-art ARMs in accuracy. Meanwhile, it also maintains fast inference because it only requires three denoising steps. These results indicate that MDiff4STR establishes a novel STR paradigm, simultaneously achieving superior accuracy and high efficiency. The contributions of this paper are threefold:

- To the best of our knowledge, we introduce MDMs into the STR task for the first time, pioneering a novel alternative to ARMs that enables a flexible trade-off between accuracy and efficiency.
- We identify two key challenges in applying the the MDM in STR: noising gap between training and inference, and overconfident predictions during inference. To address these issues, we propose six tailored noise strategies as well as a token-replacement noise mechanism.
- We present MDiff4STR, a MDM-based framework for STR that outperforms state-of-the-art ARMs in recogni-




Figure 2: The network of MDiff4STR. F_v and T_m denote the visual features and the noised tokens, respectively.

tion accuracy while achieving $3\times$ faster inference, establishing a novel paradigm for the task.

2 Related Work

Scene Text Recognition (STR), as a typical vision-language task, often heavily relies on linguistic context to achieve accurate recognition. Most efforts (Shi et al. 2016, 2019; Sheng, Chen, and Xu 2019; Li et al. 2019; Yue et al. 2020; Jiang et al. 2023; Xie et al. 2022; Zheng et al. 2024; Xu et al. 2024; Yang et al. 2024; Zhou et al. 2024; Du et al. 2025c,a; Su et al. 2025) integrate language modeling capabilities into STR using autoregressive models (ARMs). As illustrated in Fig. 1(a), ARMs predict characters through iterative decoding, explicitly modeling the contextual dependencies between characters. However, the inherently sequential nature of ARMs limits their decoding efficiency. To overcome this limitation and improve inference speed, researchers have proposed parallel decoding models (PDMs) (Wang et al. 2020, 2021; Wang, Da, and Yao 2022; Du et al. 2025b; Zhang et al. 2023; Qiao et al. 2021; Yang, Qiao, and Zhou 2025), as illustrated in Fig. 1(b). These models abandon token-to-token dependency modeling and instead generate the entire character sequence simultaneously, resulting in a significant boost in decoding speed. However, due to the absence of contextual modeling, their recognition accuracy generally lags behind that of ARMs. To strike a better balance between speed and accuracy, some works (Yu et al. 2020; Fang et al. 2021; Bautista and Atienza 2022; Na, Kim, and Park 2022; Wei et al. 2024) have introduced BERT-like architectures (Devlin et al. 2019), as depicted in Fig. 1(c). These models first produce an initial results in parallel and then refine the predictions by incorporating contextual information. While this strategy mitigates the lack of contextual understanding in purely parallel models, it can be sensitive to initial prediction errors, which may propagate during refinement (Jiang et al. 2023; Du et al. 2025d), limiting their ability to outperform ARMs. MDM, distinct from them, learns to reconstruct character sequences from partially masked inputs through a denoising process, capturing more flexible and comprehensive omnidirectional dependencies than ARMs and refinement models. Moreover, its denoising process is highly efficient and controllable, enabling accurate predictions to be produced in just a few steps, thereby achieving faster inference.

3 Method

Our MDiff4STR is a novel STR method based on the recent MDM. Fig. 2 provides details of the model architec-

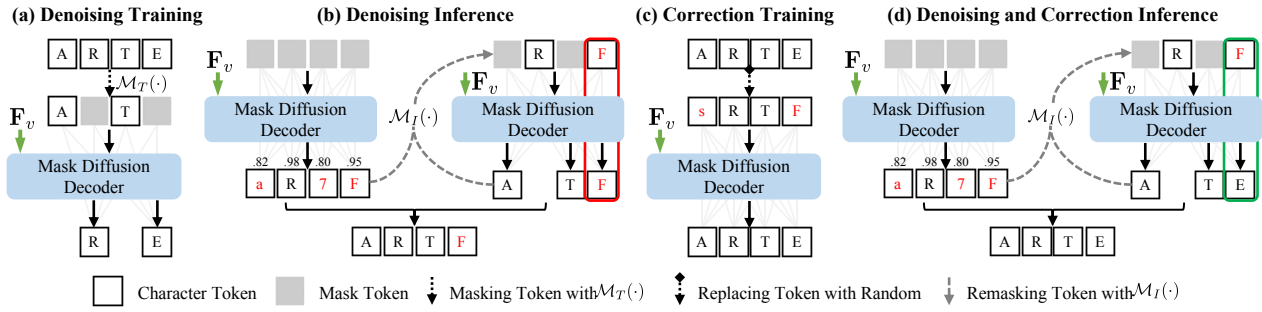


Figure 3: (a) and (b) denote the denoising training and inference of the vanilla MDM, respectively. (c) depicts the error-correction training enabled by our proposed token-replacement noise mechanism. MDiff4STR jointly leverages (a) denoising and (c) error-correction training to achieve (d) the denoising process augmented with corrective capability. Red boxes indicate errors caused by overconfident predictions, whereas green boxes highlight correct reasoning performed by MDiff4STR. \mathcal{M}_T and \mathcal{M}_I denote the mask strategy for training and the remask strategy for inference, respectively.

ture, while Fig. 3 illustrates the overall training and inference pipeline. Given an input text image $\mathbf{X} \in \mathbb{R}^{H \times W \times 3}$ and its corresponding text label $\mathbf{Y} = \{y_1, y_2, \dots, y_L\} \in \mathbb{V}^L$, where \mathbb{V} denotes the vocabulary, we use SVTRv2 (Du et al. 2025d), a visual encoder specially designed for STR, to extract image features $\mathbf{F}_v \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{4} \times D}$. On the textual side, the character sequence \mathbf{Y} is noised by the mask token to produce noised token sequence $\mathbf{T}_m \in \mathbb{R}^{L \times D}$. The mask diffusion decoder then performs a denoising process conditioned on the visual features to predict the final recognition.

3.1 Vanilla Mask Diffusion Model

In the denoising training of the vanilla MDM, the character sequence \mathbf{Y} is partially and randomly masked with a mask token [MASK], resulting in a noised version $\mathbf{Y}_m = \mathcal{M}_T(\mathbf{Y}) \in \mathbb{V}^L$. \mathcal{M}_T means the mask strategy for training. The mask diffusion decoder then recovers the original sequence \mathbf{Y} from the noised sequence \mathbf{Y}_m , conditioned on the image features \mathbf{F}_v . The entire training process can be formally described as:

$$\begin{aligned} \mathbf{F}_v &= \text{SVTRv2}(\mathbf{X}), \mathbf{T}_m = \text{Embedding}(\mathbf{Y}_m) \\ \tilde{\mathbf{T}} &= \text{MDiffDecoder}(\mathbf{F}_v, \mathbf{T}_m), \tilde{\mathbf{Y}} = \text{Classifier}(\tilde{\mathbf{T}}) \end{aligned}$$

where $\text{Embedding}(\cdot)$ is a learnable character embedding layer that maps characters into a vector space, and $\text{Classifier}(\cdot)$ maps the decoded tokens back to character.

The inference process is modeled as K -step denoising diffusion process. The first step, the token sequence is entirely set to mask tokens, representing a completely unknown character sequence. Then, $K - 1$ remask steps are applied to progressively generate the final prediction $\tilde{\mathbf{Y}}^K$:

$$\begin{aligned} \mathbf{Y}_m^1 &= [\text{MASK}]^{\otimes L}, \mathbf{T}_m^i = \text{Embedding}(\mathbf{Y}_m^i) \\ \tilde{\mathbf{T}}^i &= \text{MDiffDecoder}(\mathbf{F}_v, \mathbf{T}_m^i), \tilde{\mathbf{Y}}^i = \text{Classifier}(\tilde{\mathbf{T}}^i) \\ \mathbf{Y}_m^{i+1} &= \mathcal{M}_I(\tilde{\mathbf{Y}}^i) \end{aligned}$$

Depending on the remask strategy for inference \mathcal{M}_I , the MDM can flexibly implement multiple decoding paradigms. Specifically, as shown in Fig. 4(b), when the full

mask is adopted, parallel decoding (MDiff-PD) is formed. Fig. 4(c/d) respectively illustrate the forward and backward auto-regressive denoising processes. To maintain consistency with previous ARMs, we only consider the denoising behaviour in Fig. 4(c) as auto-regressive decoding (MDiff-AR). Fig. 4(e) corresponds to BERT-like refinement decoding (MDiff-Re). In addition, the MDM introduces a unique and efficient decoding method: confidence-guided remask. Fig. 4(f) demonstrates low-confidence remask (MDiff-LC), where after each denoising step, low-confidence tokens denoting below the average confidence score are remasked based on predicted confidence and fed into the next iteration. This approach leverages the MDM’s ability to repair uncertain predictions. However, this strategy is prone to the confidence trap, where certain tokens are repeatedly remasked in each iteration, leading to generation stagnation. To avoid this issue, MDiff-BLC (Fig. 4(g)) adopts remask for low-confidence tokens within a fixed-size local block, effectively avoiding the confidence trap. Here, the block size is set to $\frac{L}{K}$.

3.2 Training Mask Strategies (\mathcal{M}_T)

During training, as shown in Fig. 4(a), the vanilla MDM typically adopts random mask strategies to corrupt input sequences. In contrast, the inference process begins with a fully masked token sequence (Fig. 4(b)) and then is followed by a multi-step denoising process using the remask strategies \mathcal{M}_I in Fig. 4(c/d/e/f/g).

The full mask and remask strategies for inference are included in the training of the randomized mask strategy, but with minimal frequency. This creates noising gap between training and inference, leading to poor generalization. To address this issue, we use seven mask strategies (noted as \mathcal{M}_T), i.e. Fig. 4(a/b/c/d/e/f/g), from which one is uniformly sampled during training for each input sequence. This design significantly improves the model’s robustness and adaptability to the remask patterns encountered during inference.

3.3 Token-Replacement Noise Mechanism

MDMs tend to be overconfident and often assign high confidence score to incorrect predictions. For example, the char-

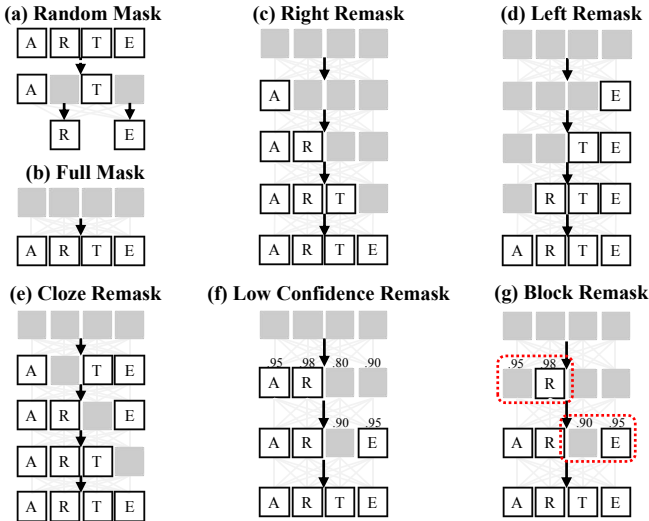


Figure 4: (a) illustrates random token masking for training. (b) shows the full mask strategy, which is used as the initial denoising step during inference. Subfigures (b–g) present various remasking strategies for inference and also serve as noise strategies in training to eliminate the noising gap.

acter “F” in Fig. 3(b) is predicted with a high confidence score of 0.95, preventing it from being remasked. As a result, the incorrect prediction cannot be corrected in subsequent steps, ultimately leading to a misrecognized output.

To address this issue, we propose a token-replacement noise mechanism. Specifically, we randomly replace certain characters in the original sequence \mathbf{Y} with another character to construct a corrupted sequence $\mathbf{Y}_r \in \mathbb{V}^L$. This sequence is then embedded into token representations $\mathbf{T}_r = \text{Embedding}(\mathbf{Y}_r) \in \mathbb{R}^{L \times D}$, which serve as input to simulate erroneous predictions. An example of such corruption is illustrated in Fig. 3(c), demonstrating how the model handles incorrect tokens during training. This method closely emulates the “erroneous yet high-confidence” scenarios encountered during inference. Subsequently, the model is trained to correct these incorrect tokens as part of the denoising task:

$$\tilde{\mathbf{T}} = \text{MDiffDecoder}(\mathbf{F}_v, \mathbf{T}_r), \tilde{\mathbf{Y}} = \text{Classifier}(\tilde{\mathbf{T}})$$

As shown in Fig. 3(d), our token-replacement noise mechanism enables robust error correction during iterative decoding. Consequently, even if the error token “F” is not remasked in a given iteration, the model can still rectify it in subsequent iterations. Additionally, the token-replacement noise mechanism presents a novel and effective noise type distinct from simple masking, suggesting that MDM noise paradigms can extend beyond masking to unlock new research avenues.

3.4 Training Objectives

MDiff4STR integrates two distinct training objectives, corresponding to the structures shown in Fig. 3(a) and 3(c), respectively. They are defined as follows:

$$\begin{aligned} \mathcal{L}_{\text{denoising}} &= -\frac{1}{l_1} \sum_{i=1}^L \mathbf{1}[\mathbf{Y}_{l_1}^i = \mathbf{M}] \log p_{\theta}(\mathbf{Y}^i | \mathbf{Y}_{l_1}) \\ \mathcal{L}_{\text{correction}} &= -\frac{1}{L} \sum_{i=1}^L \log p_{\theta}(\mathbf{Y}^i | \mathbf{Y}_{l_2}) \\ \mathcal{L}_{\text{total}} &= \mathcal{L}_{\text{denoising}} + \mathcal{L}_{\text{correction}} \end{aligned}$$

Here, $\mathbf{Y}_{l_1} = \mathbf{Y}_m$ and $\mathbf{Y}_{l_2} = \mathbf{Y}_r$. The terms l_1 and l_2 denote the number of tokens that are masked and the number of tokens that are randomly replaced with other characters, respectively. Both l_1 and l_2 are sampled uniformly from the range $[0, L]$, where L is the length of the character sequence.

In the denoising loss $\mathcal{L}_{\text{denoise}}$, $\mathbf{1}[\mathbf{Y}_{l_1}^i = \mathbf{M}]$ is an indicator function that equals 1 only if the token at position i is a [MASK] token. This ensures that the model is supervised exclusively on the masked positions. The term $p_{\theta}(\mathbf{Y}^i | \mathbf{Y}_{l_1})$ represents the probability of the model predicting the original token \mathbf{Y}^i at position i , given the masked input sequence \mathbf{Y}_{l_1} . Conversely, for the correction loss $\mathcal{L}_{\text{correction}}$, the supervision is applied across the entire sequence. This is critical because, during inference, the model has no prior knowledge of which tokens are incorrect. Therefore, the correction training requires the model to make predictions for all tokens, compelling it to learn error correction capabilities under conditions of unknown perturbations. The final training objective $\mathcal{L}_{\text{total}}$ is the sum of these two losses.

4 Experiments

4.1 Datasets and Implementation Details

For English recognition, we train MDiff4STR on *U14M-Filter* (Du et al. 2025d), which without data leakage. Then, we evaluate MDiff4STR across multiple benchmarks covering diverse scenarios. They are: 1) six common regular and irregular benchmarks (*Com*), including ICDAR 2013 (*IC13*) (KaratzasAU et al. 2013), Street View Text (*SVT*) (Wang, Babenko, and Belongie 2011), IIIT5K-Words (*IIIT5K*) (Mishra, Karteek, and Jawahar 2012), ICDAR 2015 (*IC15*) (Karatzas et al. 2015), Street View Text-Perspective (*SVTP*) (Phan et al. 2013) and *CUTE80* (Anhar et al. 2014). For IC13 and IC15, we use the versions with 857 and 1811 images, respectively; 2) Union14M-Benchmark (*U14M*) (Jiang et al. 2023), which includes seven challenging subsets: *Curve*, *Multi-Oriented*, *Artistic*, *Contextless*, *Salient*, *Multi-Words* and *General*; 3) occluded scene text dataset (*OST*) (Wang et al. 2021), which requires contextual reasoning for accurate recognition.

For Chinese recognition, we use *BCTR* (Chen et al. 2021) including four scenarios. We trained the model on integration of the four subsets and then evaluated it on the test sets: *Scene*, *Web*, *Document* (*Doc*) and *Hand-Writing* (*HW*).

We use AdamW optimizer (Loshchilov and Hutter 2019) with a weight decay of 0.05 for training. The learning rate (LR) is set to 5×10^{-4} and batchsize is set to 1024. One cycle LR scheduler (I. Loshchilov and Hutter 2017) with 1.5/4.5 epochs linear warm-up is used in all the 40/100 epochs for English and Chinese model, respectively. Word accuracy is used as the evaluation metric. Data augmentation like rotation, perspective distortion, motion blur, and gaussian noise,





	Groundtruth	V-MDiff4STR-LC	V-MDiff4STR-BLC	MDiff4STR-LC	MDiff4STR-BLC
	F A R W E S T F U N G I	Step 1: F A N <u>V</u> E S I F U N G G	F A N <u>V</u> E S I F U N G G	F A N <u>V</u> E S I F U N G I	F A N <u>V</u> E S I F U N G I
	V-MDiff4STR-PD	Score: 69.84.55.85.89.59.67.94.60.58.75.64	69.84.55.85.89.59.67.94	60.83.57.87.79.63.43.92.53.45.71.92	60.83.57.87.79.63.43.92
	Remask: * A * <u>V</u> E * * F * * G * * A * <u>V</u> E * * F * * * * * A * <u>V</u> E * * F * * G I * A * <u>V</u> E * * F * * * * *	MDiff4STR-PD	Step 2: F A R <u>V</u> E S T F U N G I	F A R <u>V</u> E S T F U N G I	F A R W E S T F U N G I
F A N V E S I F U N G I	Score: 90.84.59.85.89.81.84.94.84.89.75.84	.86.90.76.87.82.86.88.93.82.74.73.92.74.72.71.96	.77.85.82.92		
ARM _{base}	Step 3: F A * <u>V</u> E * T F U U * I	F A R <u>V</u> E S T F U U * I	F A R W E * * F * * * I	F A R W E S T F * N * I	
F A R V E S T F U N G I	Step 3: F A R <u>V</u> E S T F U N G I	F A R <u>V</u> E S T F U N G I	F A R W E S T F U N G I	F A R W E S T F U N G I	
	Groundtruth	V-MDiff4STR-LC	V-MDiff4STR-BLC	MDiff4STR-LC	MDiff4STR-BLC
	A B B I G L I A M E N T O	Step 1: I B B I G I I A M E T T O	I B B I G I I A M E T T O	B B B I G I I A M E T T O	B B B I G I I A M E T T O
	V-MDiff4STR-PD	Score: 66.98.92.58.55.89	50.91.95.96.60.98.77	66.98.92.58.55.89.50.91	92.98.84.48.75.44.50.51.97.96.72.98.67
	Remask: * B B * * I * A M E * T * * B B * * I * A * * * * * B B B * G * * M E * T * B B B * G * * * * * * * * *	MDiff4STR-PD	Step 2: A B B I G I I A M E T T O	A B B I G I I A M E T T O	A B B I G I I A M E N T O
B B B I G I I A M E T T O	Score: 73.98.92.80.78.89	80.91.95.96.72.98.83	95.95.72.99.83	94.99.91.88.88.87	55.61.97.96.82.98.67
ARM _{base}	Step 3: I B B I G I I A M E T T O	A B B I G I I A M E T T O	A B B I G I I A M E T T O	A B B I G L I A M E T T O	
I B B I G L I A M E N T O	Groundtruth: DAYTIME	ARM _{base} : CAYTIME	ReM _{base} : CAYTIME	MDiff4STR-BLC: DAYTIME	
	Groundtruth: RANCIA	ARM _{base} : BANCIA	ReM _{base} : KANCIA	MDiff4STR-BLC: RANCIA	
		Groundtruth: DUGOUT	ARM _{base} : DUGSUT	ReM _{base} : DUG_UT	MDiff4STR-BLC: DUGOUT

Figure 5: The first two figures present the denoising process of MDiff4STR, while the last three demonstrate its reasoning advantage over ARM and ReM in omnidirectional contextual modeling involving occluded or artistic text recognition. V-MDiff4STR represents indicates that the token-replacement noise mechanism is not used during training. Red characters and underline denote the misrecognition and misspelling, respectively. Red boxes indicate errors caused by overconfident predictions, whereas green boxes highlight correct reasoning performed by MDiff4STR. * indicates tokens that are remasked as the [mask].

		Com		U14M		OST		Time
		LC	BLC	LC	BLC	LC	BLC	
K	1	96.88	96.88	86.69	86.69	81.31	81.31	10.52
	2	97.19	97.19	88.05	88.05	83.69	83.69	15.56
	3	97.29	97.30	88.37	88.44	84.21	84.25	19.21
	4	97.31	97.28	88.42	88.59	84.19	84.27	23.11
	5	97.31	97.28	88.42	88.50	84.19	84.42	25.70
	6	97.31	97.29	88.42	88.57	84.19	84.33	28.74
	7	97.31	97.24	88.42	88.63	84.19	84.09	30.64
	8	97.31	97.24	88.42	88.65	84.19	84.11	32.74
N	2	96.73	96.71	87.92	87.95	81.66	81.64	12.26
	4	96.92	97.07	88.04	88.10	82.30	82.33	16.07
	6	97.29	97.30	88.37	88.44	84.21	84.25	19.21
	8	97.00	97.00	87.63	87.80	82.65	82.72	23.04

Table 1: Ablation study on the number of denoising steps K and decoder layers N in MDiff4STR.

are randomly performed. The maximum text length is set to 25. The vocabulary size $|\mathcal{V}|$ is set to 94 for English and 6624 (Li et al. 2022) for Chinese. All models are trained on 4 RTX 3090 GPUs.

4.2 Ablation Study

Tab. 1 presents the impact of the number of denoising steps (K) and decoder layers (N) on both accuracy and inference speed. Based on the results, we set $K = 3$ and $N = 6$ to achieve a balance between accuracy and efficiency.

To fairly evaluate the effectiveness of MDiff4STR, we implemented three baseline models: ARM_{base} , PDM_{base} , and ReM_{base} , using the exact same model architecture shown in Fig. 2(a) and training setting present in Sec. 4.1. As shown in the first part of Tab. 2, ARM_{base} achieves superior performance compared to all previous methods (refer to Tab. 3), which supports the robustness of our baseline models.

Decoding	\mathcal{M}_T	TRN	Com	U14M	OST	Time	
PDM_{base}	-	-	95.78	83.54	74.77	10.52	
ARM_{base}	-	-	96.88	87.34	81.03	57.95	
ReM_{base}	-	-	96.05	84.91	78.98	20.11	
MDiff-PD	R	×	95.76	84.36	76.70	10.52	
MDiff-AR	R	×	96.31	85.62	77.79	66.35	
MDiff-Re	R	×	96.24	84.93	79.41	20.11	
MDiff-LC	R	×	96.43	85.33	79.41	19.21	
MDiff-BLC	R	×	96.42	85.42	79.93	19.21	
MDiff-PD	R+All	×	96.03 ^{↑0.27}	86.00 ^{↑1.64}	78.93 ^{↑2.23}	10.52	
MDiff-AR	R+All	×	96.58 ^{↑0.26}	87.13 ^{↑1.51}	79.94 ^{↑2.14}	66.35	
MDiff-Re	R+All	×	96.77 ^{↑0.53}	86.84 ^{↑1.91}	81.45 ^{↑2.04}	20.11	
MDiff-LC	R+All	×	96.95 ^{↑0.53}	86.98 ^{↑1.65}	81.86 ^{↑2.45}	19.21	
MDiff-BLC	R+All	×	96.98 ^{↑0.57}	87.09 ^{↑1.67}	81.92 ^{↑2.00}	19.21	
MDiff-PD	R+All	✓	96.88 ^{↑1.13}	86.69 ^{↑2.33}	81.31 ^{↑4.61}	10.52	
MDiff-AR	R+All	✓	97.09 ^{↑0.78}	88.63 ^{↑3.01}	82.47 ^{↑4.68}	66.35	
MDiff-Re	R+All	✓	97.12 ^{↑0.88}	88.16 ^{↑3.23}	83.90 ^{↑4.49}	20.11	
MDiff-LC	R+All	✓	97.29 ^{↑0.87}	88.37 ^{↑3.04}	84.21 ^{↑4.80}	19.21	
MDiff-BLC	R+All	✓	97.30 ^{↑0.88}	88.44 ^{↑3.02}	84.25 ^{↑4.33}	19.21	
	Base	+(b)	+(c)	+(d)	+(e)	+(f)	+(g)
U14M	85.42	↑1.04	↑1.37	↑1.41	↑1.46	↑1.53	↑1.67

Table 2: **Top:** Influence of the training mask strategies \mathcal{M}_T and the token-replacement noise mechanism (TRN). R denotes the random mask strategy, and $R+All$ denotes using all of the mask strategies in Fig. 4 for training. **Bottom:** Ablation on the six masks in Fig. 4. Here, $+(*)$ means using the mask in Fig. 4(*) gradually.

Effectiveness of the MDM in STR. The vanilla MDM (second part of Tab. 2) is trained using only a random mask strategy. During inference, we explore five different denoising inference paradigms (MDiff-PD/AR/Re/LC/BLC, detailed in Sec. 2.2). Among these, The results show that MDiff-LC and MDiff-BLC, which represent entirely new decoding strategies distinct from previous STR methods, both achieve competitive results comparing with the existing STR meth-

		Method	Common Benchmarks (<i>Com</i>)							Avg	Union14M-Benchmark (<i>U14M</i>)							Avg	<i>OST</i>	<i>Size</i>
C T C	CRNN	TPAMI (2017)	95.8	91.8	94.6	84.9	83.1	91.0	90.21	48.1	13.0	51.2	62.3	41.4	60.4	68.2	49.24	58.0	16.2	
	SVTR-B	IJCAI (2022)	98.0	97.1	97.3	88.6	90.7	95.8	94.58	76.2	44.5	67.8	78.7	75.2	77.9	77.8	71.17	69.6	18.1	
	SVTRv2-B	ICCV (2025d)	99.2	98.0	98.7	91.1	93.5	99.0	96.57	90.6	89.0	79.3	86.1	86.2	86.7	85.1	86.14	80.0	19.8	
A R M	DAN	AAAI (2020)	97.5	94.7	96.5	87.1	89.1	94.4	93.24	74.9	63.3	63.4	70.6	70.2	71.1	76.8	70.05	61.8	27.7	
	SEED	CVPR (2020)	96.5	93.2	94.2	87.5	88.7	93.4	92.24	69.1	80.9	56.9	63.9	73.4	61.3	76.5	68.87	62.6	24.0	
	AutoSTR	ECCV (2020)	96.8	92.4	95.7	86.6	88.2	93.4	92.19	72.1	81.7	56.7	64.8	75.4	64.0	75.9	70.09	61.5	6.0	
	RoScanner	ECCV (2020)	98.5	95.8	97.7	88.2	90.1	97.6	94.65	79.4	68.1	70.5	79.6	71.6	82.5	80.8	76.08	68.6	48.0	
	PARSeq	ECCV (2022)	98.9	98.1	98.4	90.1	94.3	98.6	96.40	87.6	88.8	76.5	83.4	84.4	84.3	84.9	84.26	79.9	23.8	
	MAERec	ICCV (2023)	99.2	97.8	98.2	90.4	94.3	98.3	96.36	89.1	87.1	79.0	84.2	86.3	85.9	84.6	85.17	76.4	35.7	
	LISTER	ICCV (2023)	98.8	97.5	98.6	90.0	94.4	96.9	96.03	78.7	68.8	73.7	81.6	74.8	82.4	83.5	77.64	77.1	51.1	
	CDistNet	IICV (2024)	98.7	97.1	97.8	89.6	93.5	96.9	95.59	81.7	77.1	72.6	78.2	79.9	79.7	81.1	78.62	71.8	43.3	
	CAM	PR (2024)	98.2	96.1	96.6	89.0	93.5	96.2	94.94	85.4	89.0	72.0	75.4	84.0	74.8	83.1	80.52	74.2	58.7	
	OTE	CVPR (2024)	98.6	96.6	98.0	90.1	94.0	97.2	95.74	86.0	75.8	74.6	74.7	81.0	65.3	82.3	77.09	77.8	20.3	
	SMTR	AAAI (2025a)	99.0	97.4	98.3	90.1	92.7	97.9	95.90	89.1	87.7	76.8	83.9	84.6	89.3	83.7	85.00	73.5	15.8	
	IGTR	TPAMI (2025c)	98.7	98.4	98.1	90.5	94.9	98.3	96.48	90.4	91.2	77.0	82.4	84.7	84.0	84.4	84.86	76.3	24.1	
	MDiff4STR-S-AR		99.2	98.1	98.6	90.4	94.9	98.3	96.58	91.8	94.2	79.8	85.4	88.1	86.9	85.9	87.42	79.3	18.9	
	MDiff4STR-B-AR		99.1	98.1	99.2	91.6	96.0	98.6	97.09	93.6	94.6	81.9	86.5	88.1	89.0	86.7	88.63	82.5	31.9	
P D M	VisionLAN	ICCV (2021)	98.2	95.8	97.1	88.6	91.2	96.2	94.50	79.6	71.4	67.9	73.7	76.1	73.9	79.1	74.53	66.4	32.9	
	MGP-STR	ECCV (2022)	97.9	97.8	97.1	89.6	95.2	96.9	95.75	85.2	83.7	72.6	75.1	79.8	71.1	83.1	78.65	78.7	148	
	LPV	IJCAI (2023)	98.6	97.8	98.1	89.8	93.6	97.6	95.93	86.2	78.7	75.8	80.2	82.9	81.6	82.9	81.20	77.7	30.5	
	CPPD	TPAMI (2025b)	99.0	97.8	98.2	90.4	94.0	99.0	96.40	86.2	78.7	76.5	82.9	83.5	81.9	83.5	81.91	79.6	27.0	
		MDiff4STR-S-PD		98.6	98.1	98.5	89.1	93.8	97.9	96.00	90.4	92.3	77.9	83.2	86.0	80.5	84.6	84.98	77.4	18.9
		MDiff4STR-B-PD		98.9	98.3	98.9	90.8	95.7	98.6	96.88	92.6	93.6	79.0	84.7	86.7	84.5	85.7	86.69	81.3	31.9
R e M	SRN	CVPR (2020)	97.2	96.3	97.5	87.9	90.9	96.9	94.45	78.1	63.2	66.3	65.3	71.4	58.3	76.5	68.43	64.6	51.7	
	ABINet	CVPR (2021)	98.5	98.1	97.7	90.1	94.1	96.5	95.83	80.4	69.0	71.7	74.7	77.6	76.8	79.8	75.72	75.0	36.9	
	MATRn	ECCV (2022)	98.8	98.3	97.9	90.3	95.2	97.2	96.29	82.2	73.0	73.4	76.9	79.4	77.4	81.0	77.62	77.8	44.3	
	BUSNet	AAAI (2024)	98.3	98.1	97.8	90.2	95.3	96.5	96.06	83.0	82.3	70.8	77.9	78.8	71.2	82.6	78.10	78.7	32.1	
		MDiff4STR-S-Re		99.0	98.1	98.0	90.0	94.3	97.9	96.22	92.0	93.9	80.4	84.1	87.1	85.2	85.8	86.93	81.7	18.9
		MDiff4STR-B-Re		99.2	98.3	98.6	91.3	96.7	98.6	97.12	93.4	94.3	82.2	86.0	87.7	86.8	86.7	88.16	83.7	31.9
	MDiff4STR-S-LC		99.0	98.3	98.4	90.2	94.9	97.9	96.44	91.8	94.1	80.1	85.1	87.3	84.3	85.9	86.95	81.3	18.9	
	MDiff4STR-B-LC		99.2	98.3	99.1	91.6	97.1	98.6	97.29	93.7	94.4	82.0	86.0	87.9	87.7	86.8	88.37	84.2	31.9	
	MDiff4STR-S-BLC		99.0	98.3	98.4	90.2	94.9	97.6	96.38	91.8	94.0	80.2	85.1	87.3	84.8	85.9	87.03	81.4	18.9	
	MDiff4STR-B-BLC		99.2	98.3	99.1	91.6	97.1	98.6	97.30	93.7	94.4	82.1	86.1	87.7	88.3	86.8	88.44	84.3	31.9	

 Table 3: All the models are trained on *U14M-Filter* from scratch. *Size* denotes the number of parameters of the model ($\times 10^6$).

ods in terms of accuracy (see Tab. 3), despite trailing behind the ARM_{base} . Nonetheless, MDiff-BLC delivers a substantial advantage in efficiency, achieving an average inference speed $3\times$ faster than ARM_{base} . These findings demonstrate the potential of the MDM for the STR task.

Effectiveness of Training Mask Strategies. As shown in the third part of Tab. 2, incorporating six training mask strategies in Fig. 4 with the random mask strategy, all five denoising paradigms of MDiff4STR exhibit significant performance gains. On the *Com*, *U14M*, and *OST* the average improvements are 0.43%, 1.68%, and 2.17%, respectively. The *U14M* and *OST*, the more challenging test sets, benefit most from the six training mask strategies, which effectively bridge the noising gap between training and inference. Furthermore, as shown in the bottom part of Tab. 2, all masks contribute positively when added gradually. Among them, the full mask strategy is the most effective, which can be explained by its role as the initial denoising step, laying a critical foundation for all subsequent denoising stages.

Effectiveness of Token-Replacement Noise. As shown in the fourth part of Tab. 2, the token-replacement noise mechanism (TRN) further improves performance, boosting av-

erage accuracy over vanilla MDM by 0.91%, 2.93%, and 4.58% on *Com*, *U14M*, and *OST*, respectively. This allows MDiff4STR-LC and MDiff4STR-BLC to outperform the ARM_{base} using only three denoising steps. In addition, as shown in Fig. 5, when the model encounters high-confidence score but incorrect predictions (e.g., erroneous tokens that were not re-masked), MDiff4STR can recognize and correct these errors in subsequent denoising steps. This excellent error correction ability and denoising stability is precisely benefiting from our proposed TRN.

Effectiveness of the MDM’s Omnidirectional Language Modeling. As a novel paradigm, MDiff4STR shows its most representative strength on *OST*, where its accuracy improves from 81.03% (ARM_{base}) to 84.25%, a 3.22% gain, and a significant 5.27% increase over the ReM_{base} (78.98%). The *OST*, which focuses on occluded scenes, is designed to evaluate a model’s contextual modeling ability to infer missing or obstructed characters. Furthermore, the last three samples in Fig. 5 also illustrate that MDiff4STR outperforms ARM and ReM when encountering occluded and artistic text that requires contextual reasoning for correct recognition. These results indicate that MDM’s omnidirectional

Method	Common Benchmarks (<i>Com</i>)					Avg	<i>OST</i>	
E ² STR (2024c)	99.2	98.6	98.7	93.8	96.7	99.3	97.71	80.7
VL-Reader (2024)	99.6	99.1	98.7	92.6	97.5	99.3	97.80	86.2
CLIP4STR (2024b)	99.4	98.6	98.3	90.8	97.8	99.0	97.32	82.8
DPTR (2024a)	99.5	99.2	98.5	91.8	97.1	98.6	97.45	-
IGTR (2025c)	99.2	98.3	98.8	92.0	96.8	99.0	97.34	86.5
SVTRv2-B (2025d)	99.2	98.6	98.8	93.8	97.2	99.4	97.83	86.9
MDiff4STR-BLC	99.5	98.5	98.9	94.1	97.4	99.7	98.02	87.4

Table 4: Quantitative comparison of MDiff4STR-B with the advanced methods experienced large-scale pretraining.

contextual modeling (Shi et al. 2024; Sahoo et al. 2024) offers a clear advantage over traditional unidirectional autoregressive methods (e.g., ARM) and bidirectional refinement models (e.g., ReM).

4.3 Comparison with State-of-the-arts

To facilitate a systematic comparison, we categorize existing STR methods into four decoding paradigms: Connectionist temporal classification (CTC)-based model (Graves et al. 2006), ARM, PDM, and ReM, as summarized in Tab. 3. To eliminate the influence of model size, we report results for two MDiff4STR variants: MDiff4STR-S (18.9M parameters) and MDiff4STR-B (31.9M parameters). Under the ARM, PDM, and ReM paradigms, MDiff4STR-S and MDiff4STR-B consistently outperform competing models of similar size on *Com*, *U14M*, and *OST*. Notably, the one-step denoising version, MDiff4STR-B-PD achieves state-of-the-art results, outperforming previous bests by 0.31%, 0.55%, and 1.30% on the three benchmarks, respectively. Further improvements are observed with the dedicated decoding strategy MDiff4STR-B-BLC, which outperforms the previous best results by 0.73%, 2.30%, and 4.30% on *Com*, *U14M*, and *OST*, respectively. In particular, the 4.30% gain on *OST* demonstrates the effectiveness of MDiff4STR’s omnidirectional language modeling in capturing complex contextual dependencies, offering a novel and powerful decoding paradigm for STR.

To explore the potential of MDiff4STR in leveraging large-scale pretraining, we first pretrain it on synthetic datasets (Gupta, Vedaldi, and Zisserman 2016; Jaderberg et al. 2014) and then fine-tune it on the real-world dataset *U14M-Filter*. As shown in Tab. 4, MDiff4STR clearly benefits from pretraining: compared to training from scratch (in Tab. 3), accuracy improves by 0.72% on *Com* and a significant 3.10% on *OST*. Compared with other state-of-the-art pretrained models, MDiff4STR achieves the highest average accuracy on both *Com* (98.02%) and *OST* (87.4%), further validating the strength of its omnidirectional language modeling framework for robust and generalizable STR.

In Tab. 5, we present the results of MDiff4STR on BCTR (Chen et al. 2021), a challenging Chinese text recognition benchmark. Compared to English, Chinese text poses greater complexity due to intricate stroke structures and a significantly larger character set. Despite these challenges, MDiff4STR outperforms the previous state-of-the-art by 1.3%, 1.5%, and 0.1% on the *Scene*, *Web*, and *Document*

Method	<i>Scene</i>	<i>Web</i>	<i>Doc</i>	<i>HW</i>	Avg	<i>Size</i>
CRNN (2017)	63.8	68.2	97.0	46.1	68.76	19.5
SVTR-B (2022)	77.9	78.7	99.2	62.1	79.49	19.8
DCTC (2024)	73.9	68.5	99.4	51.0	73.20	40.8
SVTRv2-B (2025d)	83.5	83.3	99.5	67.0	83.31	22.5
ASTER (2019)	61.3	51.7	96.2	37.0	61.55	27.2
MORAN (2019)	54.6	31.5	86.1	16.2	47.10	28.5
SAR (2019)	59.7	58.0	95.7	36.5	62.48	27.8
SEED (2020)	44.7	28.1	91.4	21.0	46.30	36.1
MASTER (2021)	62.8	52.1	84.4	26.9	56.55	62.8
TransOCR (2021)	71.3	64.8	97.1	53.0	71.55	83.9
PARSeq (2022)	84.2	82.8	99.5	63.0	82.37	28.9
CCR-CLIP (2023)	71.3	69.2	98.3	60.3	74.78	62.0
CAM (2024)	76.0	69.3	98.1	59.2	76.80	135
MAERec (2023)	84.4	83.0	99.5	65.6	83.13	40.8
LISTER (2023)	79.4	79.5	99.2	58.0	79.02	55.0
DPTR (2024a)	80.0	79.6	98.9	64.4	80.73	68.0
IGTR-AR (2025c)	82.0	81.7	99.5	63.8	81.74	29.2
SMTR (2025a)	83.4	83.0	99.3	65.1	82.68	20.8
MDiff4STR-S-AR	84.4	83.8	99.5	67.1	83.71	23.9
MDiff4STR-B-AR	85.1	84.2	99.6	67.3	84.04	36.9
CPPD (2025b)	82.7	82.4	99.4	62.3	81.72	32.1
MDiff4STR-S-PD	82.4	82.5	99.4	60.9	81.31	23.9
MDiff4STR-B-PD	83.4	82.2	99.5	61.8	81.96	36.9
ABINet (2021)	66.6	63.2	98.2	53.1	70.28	53.1
MDiff4STR-B-Re	84.6	83.9	99.5	65.4	83.39	23.9
MDiff4STR-B-Re	85.6	84.7	99.6	67.0	84.23	36.9
MDiff4STR-S-LC	85.2	84.1	99.6	66.0	83.72	23.9
MDiff4STR-B-LC	85.6	84.8	99.6	66.5	84.11	36.9
MDiff4STR-S-BLC	85.2	84.1	99.6	66.7	83.89	23.9
MDiff4STR-B-BLC	85.7	84.7	99.6	67.0	84.25	36.9

Table 5: Results on Chinese text dataset.

(*Doc*) subsets, respectively, and achieves comparable performance on the *Hand-Writing (HW)* subset. These results highlight the robustness and adaptability of MDiff4STR to non-Latin scripts, further validating its generalizability in multilingual text recognition tasks.

5 Conclusion

In this paper, we introduced the MDM into the STR task and proposed a novel STR framework, MDiff4STR. Leveraging MDM’s omnidirectional language modeling capability, MDiff4STR surpasses the previous prominent ARMs and BERT-like models. To overcome two key challenges in applying the MDM to STR, the noising gap between training and inference and the overconfidence in predictions during inference, we proposed the training noise strategy that aligns with inference behavior and the token-replacement noise mechanism. Extensive experiments demonstrate that MDiff4STR can flexibly support multiple decoding paradigms and achieves state-of-the-art results across a wide range of challenging scenarios, including regular, irregular, occluded, and Chinese text, with or without the use of pretraining. Furthermore, its dedicated low-confidence denoising inference surpasses ARMs with only three denoising steps, establishing a novel and effective paradigm for STR. These findings highlight the potential of MDiff4STR to advance the STR field and offer valuable insights for future research.

Acknowledgments

This work was supported by National Natural Science Foundation of China (Nos. 62427819, 62576026)

References

- Anhar, R.; Palaiahnakote, S.; Chan, C. S.; and Tan, C. L. 2014. A robust arbitrary text detection system for natural scene images. *Expert Syst. Appl.*, 41(18): 8027–8048.
- Bautista, D.; and Atienza, R. 2022. Scene Text Recognition with Permuted Autoregressive Sequence Models. In *ECCV*, 178–196.
- Chen, J.; Li, B.; and Xue, X. 2021. Scene Text Telescope: Text-Focused Scene Image Super-Resolution. In *CVPR*, 12021–12030.
- Chen, J.; Yu, H.; Ma, J.; Guan, M.; Xu, X.; Wang, X.; Qu, S.; Li, B.; and Xue, X. 2021. Benchmarking Chinese Text Recognition: Datasets, Baselines, and an Empirical Study. *CoRR*, abs/2112.15093.
- Chen, X.; Jin, L.; Zhu, Y.; Luo, C.; and Wang, T. 2022. Text Recognition in the Wild: A Survey. *ACM Comput. Surv.*, 54(2): 42:1–42:35.
- Cheng, C.; Wang, P.; Da, C.; Zheng, Q.; and Yao, C. 2023. LISTER: Neighbor Decoding for Length-Insensitive Scene Text Recognition. In *ICCV*, 19484–19494.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*, 4171–4186.
- Du, Y.; Chen, Z.; Jia, C.; Gao, X.; and Jiang, Y.-G. 2025a. Out of Length Text Recognition with Sub-String Matching. In *AAAI*, 2798–2806.
- Du, Y.; Chen, Z.; Jia, C.; Yin, X.; Li, C.; Du, Y.; and Jiang, Y.-G. 2025b. Context Perception Parallel Decoder for Scene Text Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(6): 4668–4683.
- Du, Y.; Chen, Z.; Jia, C.; Yin, X.; Zheng, T.; Li, C.; Du, Y.; and Jiang, Y.-G. 2022. SVTR: Scene Text Recognition with a Single Visual Model. In *IJCAI*, 884–890.
- Du, Y.; Chen, Z.; Su, Y.; Jia, C.; and Jiang, Y.-G. 2025c. Instruction-Guided Scene Text Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 47(4): 2723–2738.
- Du, Y.; Chen, Z.; Xie, H.; Jia, C.; and Jiang, Y.-G. 2025d. SVTRv2: CTC Beats Encoder-Decoder Models in Scene Text Recognition. In *ICCV*, 20147–20156.
- Fang, S.; Xie, H.; Wang, Y.; Mao, Z.; and Zhang, Y. 2021. Read Like Humans: Autonomous, Bidirectional and Iterative Language Modeling for Scene Text Recognition. In *CVPR*, 7098–7107.
- Graves, A.; Fernández, S.; Gomez, F.; and Schmidhuber, J. 2006. Connectionist Temporal Classification: Labelling Unsegmented Sequence Data with Recurrent Neural Networks. In *ICML*, 369–376.
- Gupta, A.; Vedaldi, A.; and Zisserman, A. 2016. Synthetic Data for Text Localisation in Natural Images. In *CVPR*, 2315–2324.
- Jaderberg, M.; Simonyan, K.; Vedaldi, A.; and Zisserman, A. 2014. Synthetic Data and Artificial Neural Networks for Natural Scene Text Recognition. *CoRR*, abs/1406.2227.
- Jiang, Q.; Wang, J.; Peng, D.; Liu, C.; and Jin, L. 2023. Revisiting Scene Text Recognition: A Data Perspective. In *ICCV*, 20486–20497.
- Karatzas, D.; Gomez-Bigorda, L.; Nicolaou, A.; Ghosh, S.; Bagdanov, A.; Iwamura, M.; Matas, J.; Neumann, L.; Chandrasekhar, V. R.; Lu, S.; Shafait, F.; Uchida, S.; and Valveny, E. 2015. ICDAR 2015 competition on Robust Reading. In *ICDAR*, 1156–1160.
- KaratzasAU, D.; ShafaitAU, F.; UchidaAU, S.; IwamuraAU, M.; i. BigordaAU, L. G.; MestreAU, S. R.; MasAU, J.; MottaAU, D. F.; AlmazàAU, J. A.; and de las Heras, L. P. 2013. ICDAR 2013 Robust Reading Competition. In *ICDAR*, 1484–1493.
- Li, C.; Liu, W.; Guo, R.; Yin, X.; Jiang, K.; Du, Y.; Du, Y.; Zhu, L.; Lai, B.; Hu, X.; Yu, D.; and Ma, Y. 2022. PP-OCrv3: More Attempts for the Improvement of Ultra Lightweight OCR System. *CoRR*, abs/2206.03001.
- Li, H.; Wang, P.; Shen, C.; and Zhang, G. 2019. Show, attend and read: A simple and strong baseline for irregular text recognition. In *AAAI*, 8610–8617.
- Loshchilov, I.; and Hutter, F. 2019. Decoupled Weight Decay Regularization. In *ICLR*.
- Lu, N.; Yu, W.; Qi, X.; Chen, Y.; Gong, P.; Xiao, R.; and Bai, X. 2021. MASTER: Multi-aspect non-local network for scene text recognition. *Pattern Recognit.*, 117: 107980.
- Luo, C.; Jin, L.; and Sun, Z. 2019. MORAN: A Multi-Object Rectified Attention Network for Scene Text Recognition. *Pattern Recognit.*, 90: 109–118.
- Mishra, A.; Karteek, A.; and Jawahar, C. V. 2012. Scene Text Recognition using Higher Order Language Priors. In *BMVC*, 1–11.
- Na, B.; Kim, Y.; and Park, S. 2022. Multi-modal Text Recognition Networks: Interactive Enhancements Between Visual and Semantic Features. In *ECCV*, 446–463.
- Nie, S.; Zhu, F.; You, Z.; Zhang, X.; Ou, J.; Hu, J.; Zhou, J.; Lin, Y.; Wen, J.; and Li, C. 2025. Large Language Diffusion Models. *CoRR*, abs/2502.09992.
- Phan, T. Q.; Shivakumara, P.; Tian, S.; and Tan, C. L. 2013. Recognizing Text with Perspective Distortion in Natural Scenes. In *CVPR*, 569–576.
- Qiao, Z.; Zhou, Y.; Wei, J.; Wang, W.; Zhang, Y.; Jiang, N.; Wang, H.; and Wang, W. 2021. PIMNet: a parallel, iterative and mimicking network for scene text recognition. In *ACM MM*, 2046–2055.
- Qiao, Z.; Zhou, Y.; Yang, D.; Zhou, Y.; and Wang, W. 2020. SEED: Semantics Enhanced Encoder-Decoder Framework for Scene Text Recognition. In *CVPR*, 13525–13534.
- Sahoo, S. S.; Arriola, M.; Schiff, Y.; Gokaslan, A.; Marroquin, E.; Chiu, J. T.; Rush, A.; and Kuleshov, V. 2024. Simple and Effective Masked Diffusion Language Models. In *NeurIPS*.

- Sheng, F.; Chen, Z.; and Xu, B. 2019. NRTR: A No-Recurrence Sequence-to-Sequence Model for Scene Text Recognition. In *ICDAR*, 781–786.
- Shi, B.; Bai, X.; and Yao, C. 2017. An End-to-End Trainable Neural Network for Image-Based Sequence Recognition and Its Application to Scene Text Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(11): 2298–2304.
- Shi, B.; Wang, X.; Lyu, P.; Yao, C.; and Bai, X. 2016. Robust scene text recognition with automatic rectification. In *CVPR*, 4168–4176.
- Shi, B.; Yang, M.; Wang, X.; Lyu, P.; Yao, C.; and Bai, X. 2019. ASTER: An Attentional Scene Text Recognizer with Flexible Rectification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 41(9): 2035–2048.
- Shi, J.; Han, K.; Wang, Z.; Doucet, A.; and Titsias, M. K. 2024. Simplified and Generalized Masked Diffusion for Discrete Data. In *NeurIPS*.
- Su, Y.; Chen, Z.; Du, Y.; Wu, Z.; Xie, H.; and Jiang, Y.-G. 2025. LRANet++: Low-Rank Approximation Network for Accurate and Efficient Text Spotting. *CoRR*, abs/2511.05818.
- I. Loshchilov; and Hutter, F. 2017. SGDR: Stochastic Gradient Descent with Warm Restarts. In *ICLR*.
- Wang, K.; Babenko, B.; and Belongie, S. 2011. End-to-end scene text recognition. In *ICCV*, 1457–1464.
- Wang, P.; Da, C.; and Yao, C. 2022. Multi-Granularity Prediction for Scene Text Recognition. In *ECCV*, 339–355.
- Wang, T.; Zhu, Y.; Jin, L.; Luo, C.; Chen, X.; Wu, Y.; Wang, Q.; and Cai, M. 2020. Decoupled Attention Network for Text Recognition. In *AAAI*, 12216–12224.
- Wang, Y.; Xie, H.; Fang, S.; Wang, J.; Zhu, S.; and Zhang, Y. 2021. From Two to One: A New Scene Text Recognizer With Visual Language Modeling Network. In *ICCV*, 14194–14203.
- Wei, J.; Zhan, H.; Lu, Y.; Tu, X.; Yin, B.; Liu, C.; and Pal, U. 2024. Image as a Language: Revisiting Scene Text Recognition via Balanced, Unified and Synchronized Vision-Language Reasoning Network. In *AAAI*, 5885–5893.
- Xie, X.; Fu, L.; Zhang, Z.; Wang, Z.; and Bai, X. 2022. Toward Understanding WordArt: Corner-Guided Transformer for Scene Text Recognition. In *ECCV*, 303–321.
- Xu, J.; Wang, Y.; Xie, H.; and Zhang, Y. 2024. OTE: Exploring Accurate Scene Text Recognition Using One Token. In *CVPR*, 28327–28336.
- Yang, M.; Yang, B.; Liao, M.; Zhu, Y.; and Bai, X. 2024. Class-Aware Mask-guided feature refinement for scene text recognition. *Pattern Recognit.*, 149: 110244.
- Yang, X.; Qiao, Z.; and Zhou, Y. 2025. IPAD: Iterative, Parallel, and Diffusion-Based Network for Scene Text Recognition. *Int. J. Comput. Vis.*, 133: 5589–5609.
- You, Z.; Nie, S.; Zhang, X.; Hu, J.; Zhou, J.; Lu, Z.; Wen, J.; and Li, C. 2025. LLaDA-V: Large Language Diffusion Models with Visual Instruction Tuning. *CoRR*, abs/2505.16933.
- Yu, D.; Li, X.; Zhang, C.; Liu, T.; Han, J.; Liu, J.; and Ding, E. 2020. Towards accurate scene text recognition with semantic reasoning networks. In *CVPR*, 12113–12122.
- Yu, H.; Wang, X.; Li, B.; and Xue, X. 2023. Chinese Text Recognition with A Pre-Trained CLIP-Like Model Through Image-IDS Aligning. In *ICCV*, 11909–11918.
- Yue, X.; Kuang, Z.; Lin, C.; Sun, H.; and Zhang, W. 2020. RobustScanner: Dynamically enhancing positional clues for robust text recognition. In *ECCV*, 135–151.
- Zhang, B.; Xie, H.; Wang, Y.; Xu, J.; and Zhang, Y. 2023. Linguistic More: Taking a Further Step toward Efficient and Accurate Scene Text Recognition. In *IJCAI*, 1704–1712.
- Zhang, H.; Yao, Q.; Yang, M.; Xu, Y.; and Bai, X. 2020. AutoSTR: Efficient backbone search for scene text recognition. In *ECCV*, 751–767.
- Zhang, Z.; Lu, N.; Liao, M.; Huang, Y.; Li, C.; Wang, M.; and Peng, W. 2024. Self-Distillation Regularized Connectionist Temporal Classification Loss for Text Recognition: A Simple Yet Effective Approach. In *AAAI*, 7441–7449.
- Zhao, S.; Du, Y.; Chen, Z.; and Jiang, Y.-G. 2024a. Decoder Pre-Training with only Text for Scene Text Recognition. In *ACM MM*, 5191–5200. ISBN 9798400706868.
- Zhao, S.; Quan, R.; Zhu, L.; and Yang, Y. 2024b. CLIP4STR: A Simple Baseline for Scene Text Recognition With Pre-Trained Vision-Language Model. *IEEE Trans. Image Process.*, 33: 6893–6904.
- Zhao, Z.; Tang, J.; Lin, C.; Wu, B.; Huang, C.; Liu, H.; Tan, X.; Zhang, Z.; and Xie, Y. 2024c. Multi-modal In-Context Learning Makes an Ego-evolving Scene Text Recognizer. In *CVPR*, 15567–15576.
- Zheng, T.; Chen, Z.; Bai, J.; Xie, H.; and Jiang, Y.-G. 2023. TPS++: Attention-Enhanced Thin-Plate Spline for Scene Text Recognition. In *IJCAI*, 1777–1785.
- Zheng, T.; Chen, Z.; Fang, S.; Xie, H.; and Jiang, Y.-G. 2024. CDistNet: Perceiving multi-domain character distance for robust text recognition. *Int. J. Comput. Vis.*, 132(2): 300–318.
- Zhong, H.; Yang, Z.; Li, Z.; Wang, P.; Tang, J.; Cheng, W.; and Yao, C. 2024. VL-Reader: Vision and Language Reconstructor is an Effective Scene Text Recognizer. In *ACM MM*, 4207–4216.
- Zhou, B.; Qu, Y.; Wang, Z.; Li, Z.; Zhang, B.; and Xie, H. 2024. Focus on the Whole Character: Discriminative Character Modeling for Scene Text Recognition. In *IJCAI*, 1762–1770.