

Dereflection Any Image with Diffusion Priors and Diversified Data

Jichen Hu^{1*}, Chen Yang^{1*}, Zanwei Zhou¹, Jiemin Fang^{2†}, Qi Tian², Wei Shen^{1✉†}

¹ MoE Key Lab of Artificial Intelligence, AI Institute, School of Computer Science, Shanghai Jiao Tong University

² Huawei Inc.

{hujichen, sjtu19zzw, wei.shen}@sjtu.edu.cn,

{chenyang.res, jaminfang}@gmail.com, tian.qi1@huawei.com

Abstract

Reflection removal of a single image remains a highly challenging task due to the complex entanglement between target scenes and unwanted reflections. Despite significant progress, existing methods are hindered by the scarcity of high-quality, diverse data and insufficient restoration priors, resulting in limited generalization across various real-world scenarios. In this paper, we propose Dereflection Any Image, a comprehensive solution with an efficient data preparation pipeline and a generalizable model for robust reflection removal. First, we introduce a dataset named Diverse Reflection Removal (DRR) created by randomly rotating reflective mediums in target scenes, enabling variation of reflection angles and intensities and setting a new benchmark in scale, quality, and diversity. Second, we propose a diffusion-based framework with one-step diffusion for deterministic outputs and fast inference. To ensure stable learning, we design a three-stage progressive training strategy that includes reflection-invariant finetuning to encourage consistent outputs across varying reflection patterns that characterize our dataset. Extensive experiments show that our method achieves SOTA performance on both common benchmarks and challenging in-the-wild images, showing superior generalization across diverse real-world scenes.

Project Page — <https://abuuu122.github.io/DAI.github.io>

Introduction

Capturing images through glass or other reflective mediums often introduces unwanted reflections, which degrade the visibility of the underlying target scene, resulting in a mixed image with two layers superposed. These reflections significantly degrade both visual aesthetics and usages in downstream tasks (Liu et al. 2020; Wan et al. 2021). Developing strong reflection-removal (**dereflection** for short) methods is essential for practical applications.

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

*Equal contributions.

†Project lead.

✉Corresponding authors.

“Dereflection” may not be an accurate word, which is used to represent removing reflection for simplicity. This name is to respect (Kirillov et al. 2023; Yang et al. 2024).

Traditional methods (Li and Brown 2014; Shih et al. 2015; Wan et al. 2016) for reflection removal rely on empirical assumptions like reflections being blurred or exhibiting ghosting effects, which often fail in real-world scenarios. Learning-based methods (Hong et al. 2024; Hu and Guo 2023; Hu, Wang, and Guo 2024; Song et al. 2023; Zhong et al. 2024; Dong et al. 2021; Hu and Guo 2021; Zhu et al. 2024) attempt to learn dereflection capabilities through paired data. However, such paired data is difficult to obtain in real-world settings, *e.g.*, capturing museum exhibits with and without reflections typically requires physically removing the glass from display cases, which is often impractical. Consequently, existing real-world datasets (Wan et al. 2017; Zhang, Ng, and Chen 2018; Li et al. 2020; Zhu et al. 2024) are limited in scale and diversity, failing to capture a wide range of reflection patterns encountered in practice. Moreover, synthetic datasets (Wen et al. 2019; Kim, Huo, and Yoon 2020; Hu and Guo 2023) often suffer from significant domain gaps, as their reflection patterns and image characteristics do not fully align with real-world scenarios.

To address these limitations, we propose an efficient data collection pipeline and introduce Diverse Reflection Removal (DRR), a 4K dataset with diverse reflection patterns. Specifically, for each view of the capturing device, reflective mediums are randomly rotated with the target scene. Then we capture mixed videos and decompose them into frames to form data pairs. This approach allows us to flexibly vary reflection angles, intensities, and scene diversity, significantly enhancing the realism and variety of reflection patterns compared to existing datasets (Wan et al. 2017; Zhang, Ng, and Chen 2018; Li et al. 2020; Zhu et al. 2024), as shown in Table 1 and Fig. 2. We also construct synthetic pairs to strengthen the data sufficiency, which are filtered with the CLIP score (Radford et al. 2021) to guarantee the data realism. Besides extending the diversity of training data, we propose to leverage the generalization capabilities of diffusion models, which excel in image-to-image translation tasks even tuned with finite data (Xia et al. 2023; Tumanyan et al. 2023). Our model integrates ControlNet (Zhang, Rao, and Agrawala 2023) to use the mixed image as a conditioning signal. To ensure deterministic outputs for reliable transmission layer recovery, we build our generative prior-based model on one-step diffusion (Ye et al. 2024; Xu et al. 2024), enabling both deterministic results



Figure 1: Our model demonstrates strong and general reflection removal capabilities. Upper: Original images with reflections. Bottom: Results generated by our model. The scenarios include glass, plastic, water surfaces, *etc.*

and fast inference. To ensure stable learning of translation from complex mixed images to dereflected ones, we design a three-stage progressive training strategy. Initially, we employ foundation training with image pairs to get basic performance. We further exploit the unique characteristics of our dataset through a reflection-invariant finetuning strategy. Since our dataset contains identical transmission scenes with varying reflection patterns, we train the model to produce consistent outputs despite these variations, enhancing generalization by focusing on invariant properties of transmission scenes rather than variable reflection characteristics. Finally, a cross-latent decoder is trained to mitigate blurriness and preserve details. The new dataset, code and model checkpoints will be released.

Our key contributions are summarized as follows:

- We introduce an efficient pipeline for data collection, and present DRR, a high-quality dataset featuring diverse reflections with varying angles, fostering future advancements in the field.
- We design a novel diffusion-based framework with a progressive training strategy ensuring both stable optimization and strong generalization across diverse reflection types, as shown in Fig. 1.
- Extensive experiments demonstrate that our method achieves SOTA performance not only on benchmark datasets but also on challenging in-the-wild images captured by mobile devices, exhibiting superior generalization across diverse real-world reflection scenarios.

Related Work

Reflection Removal

Reflection removal is an ill-posed problem that aims to separate the reflection layer from a mixed image and recover the underlying transmission layer. It is well-established that additional information can significantly simplify this task. For instance, multi-view (Niklaus et al. 2021; Li, Chan, and Lun 2020; Guo, Cao, and Ma 2014) information can guide robust reflection removal, but capturing image sequences is often redundant. Flash-based methods (Lei, Jiang, and Chen 2023; Wang et al. 2024c) utilize flash/no-flash image pairs, while polarization-based techniques (Lei et al. 2020) exploit different polarization angles. However, these approaches rely

on specialized equipment, limiting their general applicability. As a result, we focus on the most challenging yet widely applicable task: single-image reflection removal (SIRR).

Traditional methods for reflection removal rely on hand-crafted priors, such as the assumption that reflections are blurred (Li and Brown 2014; Wan et al. 2016) or exhibit ghosting effects (Shih et al. 2015). However, these assumptions frequently break down in complex real-world scenarios, resulting in suboptimal performance. Universal image restoration techniques (Chen et al. 2022; Kawar et al. 2022; Wang et al. 2022) could serve as a potential solution, however, they cannot outperform methods designed specifically for single image reflection removal. In recent years, learning-based methods (Zhao et al. 2025; Hong et al. 2024; Hu and Guo 2023; Hu, Wang, and Guo 2024; Song et al. 2023; Zhong et al. 2024; Dong et al. 2021; Hu and Guo 2021; Fan et al. 2017; Wei et al. 2019; Yang et al. 2018; Wang et al. 2024b) have become the mainstream approach, enabling models to remove reflections by training on large datasets. While some real datasets (Wan et al. 2017; Zhang, Ng, and Chen 2018; Li et al. 2020; Zhu et al. 2024) have been collected, their scale and quality remain insufficient. To address this, some methods employ empirical formulations (Hu and Guo 2023), physics-based rendering (Kim, Huo, and Yoon 2020), and tailored models (Wen et al. 2019) to generate synthetic data. Nevertheless, these approaches often struggle to bridge the gap between synthetic and real-world data. In this work, we introduce a comprehensive data collection pipeline to acquire a large volume of real-world data DRR and a data filtering strategy to enhance the quality of synthetic data.

Diffusion Model

Diffusion models (Ho, Jain, and Abbeel 2020; Rombach et al. 2022; Song, Meng, and Ermon 2020) have emerged as a powerful framework for image generation and restoration, leveraging an iterative denoising process that learns the underlying data distribution. Their ability to produce high-quality, diverse outputs has garnered significant attention in recent years. Notably, diffusion models have been successfully applied to a wide range of image restoration tasks (Kawar et al. 2022), such as deblurring (Whang



Figure 2: Our dataset contains a diverse collection of scenes, each accompanied by multiple reflection images. As illustrated in the figure, the ground truth transmission layer is highlighted in red boxes, while the remaining images represent various mixed images. The dataset demonstrates remarkable diversity, encompassing **indoor**, **outdoor**, and **object-centric scenes**. All image pairs maintain high resolution with rich textual details. (Best viewed on screen.)

	SIR ²	Real	Nature	RRW	DRR (Ours)
Year	2017	2018	2020	2023	2025
Videos	×	×	×	✓	✓
Angles	×	×	×	×	✓
Usage	Test	Train/Test	Train/Test	Train	Train/Test
Pairs	500	89/20	200/20	14952	23303/400
Scenes	126	89/20	68	150	217/40
Avg. res.	540*400	1152*930	598*398	2580*1460	3840*2160

Table 1: Summary of existing real reflection datasets. Compared to these datasets, our proposed DRR dataset demonstrates significant advantages in three key aspects: (1) greater diversity in reflection (varying glass angles, reflected contents, reflection intensity) and scenes (various environment conditions), (2) a substantially larger collection of image pairs, and (3) superior image quality with higher resolution.

et al. 2022), super-resolution (Wang et al. 2024d,a), inpainting (Lugmayr et al. 2022; Xie et al. 2023) and LLIE (Zhou et al. 2025; Lan et al. 2025; Yi et al. 2023), demonstrating exceptional performance in recovering fine-grained details. These characteristics make diffusion models particularly well-suited for the challenging task of reflection removal, where disentangling complex visual structures and maintaining image fidelity are critical. Recent work, such as L-DiffER (Hong et al. 2024), utilizes ControlNet (Zhang, Rao, and Agrawala 2023) to inject information from the mixed image and iteratively denoise it. However, this approach requires multiple steps to recover the transmission

layer and text prompt to guide the process. To address this limitation, we propose an alternative solution based on a recently proposed one-step denoising strategy (Ye et al. 2024; Xu et al. 2024). It enables stable and deterministic results while significantly accelerating the process, making it more practical for real-world applications.

Method

In this section, we first present our data collection pipeline, which captures diverse real-world data and generates high-quality synthetic data. Next, we introduce our diffusion-based framework specifically designed for reflection removal. Finally, we propose a progressive training strategy to fully leverage the diversity of our dataset.

Data Collection Pipeline

Formally, we define the target scene as transmission layer \mathbf{T} , which is superposed with undesired reflection layer \mathbf{R} , resulting in the mixed image \mathbf{M} . Given that the core objective is to translate mixed image \mathbf{M} to transmission layer \mathbf{T} , we design a comprehensive data collection pipeline to acquire aligned image pairs of \mathbf{M} and \mathbf{T} , including our diverse real-world data, and synthetic data serving as a supplementary, demonstrated in Fig. 3.

Real Data Our dataset is captured using a Nikon Z50 camera and three mobile phones mounted on a fixed tripod, with a portable glass slab positioned in front of the lens. By rotating the glass at different angles within a sequence, we generate reflection images with varying reflection and intensity.

The corresponding ground-truth transmission images are acquired by removing the glass entirely. To ensure diversity and robustness, we carefully vary key parameters, including scenes (indoor, outdoor, object-centric), lighting conditions (skylight, sunlight, incandescent), glass thickness (3 mm and 8 mm), camera-to-glass distance, viewing angles, exposure values, and aperture settings.

The dataset consists of 257 unique scenes, each captured with two glass thicknesses to ensure reflection diversity. It is partitioned into a training set (217 scenes, 23,303 image pairs) and a testing set (40 scenes, 400 image pairs). All training images are captured in 4K resolution (3840×2160 pixels) to provide high visual fidelity for model training. The testing set is further divided into two subsets: DRR-S, containing standard reflections captured with the same camera used in training, and DRR-C, containing challenging reflections captured using three mobile phones. This division enables a comprehensive evaluation of model performance across varying reflection complexities and real-world scenarios.

To address spatial shifts caused by glass refraction, we employ a robust post-processing pipeline. Inspired by (Wan et al. 2017), we use Scale-Invariant Feature Transform (SIFT) (Lowe 2004) for feature point detection and Random Sample Consensus (RANSAC) for precise alignment between reflection images and their ground-truth transmission pairs.

Compared to existing datasets, our collection offers superior image quality, substantial quantity, and exceptional diversity, derived from varying glass angles, reflected contents, reflection intensity and various environment conditions. Our method has faster capturing speed and is easier to scale up. The existing RRW dataset focuses on object movement and occlusions in front of the glass, lacking the variation in angular perspectives and diverse reflection content from the real scenes that characterize our dataset. A detailed comparison is provided in Table 1.

Synthetic Data Supplementary To complement real-world datasets, we generate synthetic data to enhance training diversity. Following DSRNet (Hu and Guo 2023), we synthesize mixed images from transmission and reflection layers, then use CLIP (Radford et al. 2021) to filter for realism, retaining 20,833 high-quality pairs. Our synthetic data also supports multiple reflections per scene, further enriching the training set.

Framework

Our proposed dataset offers a comprehensive and diverse training resource, however, it is impractical to encompass all possible reflection types. Fortunately, diffusion models provide an effective solution. By leveraging the powerful generative priors of diffusion models, we can maximize the utility of limited data and address more challenging reflection scenarios. Our diffusion-based model is composed of three key parts, a U-net for one-step denoising, a ControlNet to input the mixed image, and a cross-latent decoder to preserve details, Shown in Fig. 4.

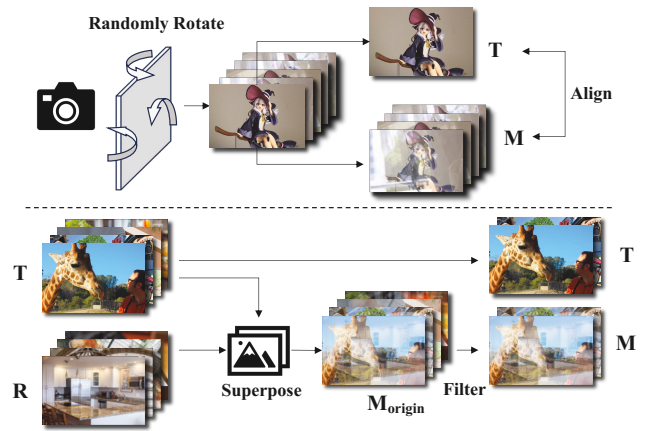


Figure 3: Data collection pipeline of real (above) and synthetic (below) data. Real data is captured by recording videos while rotating a glass panel at various angles, then processed to align mixed images with their ground-truth transmission layers. Synthetic data is generated by randomly chosen coefficients and filtered to produce high-quality image pairs.




Figure 4: Our proposed framework. It consists of a U-net with one-step denoising strategy, a ControlNet to input the mixed image processed by the encoder \mathcal{E} , and a cross-latent decoder \mathcal{D} to mitigate blurriness and preserve details.

Traditional diffusion models generate results through iterative denoising from Gaussian noise. But we use a one-step denoising strategy as an alternative, the target latent is predicted with only one step by the U-net, resulting in deterministic results and fast inference. ControlNet (Zhang, Rao, and Agrawala 2023) is initially designed to impose structural constraints on generated images. In our work, we adapt it as a mechanism to inject information from mixed images. To address the inherent limitations of one-step diffusion, such as over-smoothed results and potential shifts in details during the denoising process, we add a cross-latent decoder \mathcal{D} to preserve the high-frequency details of the input mixed image. Inspired by (Wang et al. 2024c; Ronneberger, Fischer, and Brox 2015), the multi-scale latent features extracted by the encoder are directly connected to the decoder with zero convolution, creating a shortcut that preserve the initial information of the mixed image.

Progressive Training

To effectively train our reflection removal model, we adopt a progressive training strategy that decomposes the learning process into three distinct yet interconnected phases, as illustrated in Fig. 5. This hierarchical approach enables stable optimization and ensures that each component of our model is properly trained to handle the complex task of reflection removal.

Foundation Training The initial stage focuses on establishing the fundamental reflection removal capability by jointly training the ControlNet and the upsampling blocks of the U-Net. We employ the following one-step diffusion loss function:

$$\mathcal{L}_{diff} = \mathbb{E}_{z_t, \mathbf{M}} \left[\|z_t - \mu_{\theta, \phi}^{z_t}(\mathbf{M})\|_2^2 \right], \quad (1)$$

where θ and ϕ refer to the parameters of U-Net and ControlNet, z_t is the latent representation after adding noise over t steps with $t \in (0, T)$. (in Stable Diffusion, T is set to 1,000). For notational simplicity, we leave the full formulation in the supplementary material.

Reflection-Invariant Fine-tuning Building upon the foundation, we introduce a novel reflection-invariant fine-tuning strategy that leverages the unique characteristics of our dataset, where each scene contains multiple mixed images with varying reflections. Our key insight is that the model should produce consistent results regardless of the specific reflection patterns present in the input. This is achieved by incorporating a consistent loss:

$$\mathcal{L}_{con} = \mathbb{E}_{\mathbf{M}_1, \mathbf{M}_2} \left[\|\mu_{\theta, \phi}^{z_t}(\mathbf{M}_1) - \mu_{\theta, \phi}^{z_t}(\mathbf{M}_2)\|_2^2 \right], \quad (2)$$

$$\mathcal{L} = \mathcal{L}_{diff}(\mathbf{M}_1) + \mathcal{L}_{diff}(\mathbf{M}_2) + \mathcal{L}_{con}(\mathbf{M}_1, \mathbf{M}_2). \quad (3)$$

The reflection-invariant finetuning enhances generalization by focusing on invariant properties of transmission scenes rather than variable reflection characteristics, making the model more robust and generalize to complex reflection scenarios.

Cross-Latent Decoder In the final stage, we freeze the previously trained components and focus on training the cross-latent decoder to mitigate blurriness and preserve high-frequency details. This is accomplished through a comprehensive image reconstruction loss:

$$\mathcal{L}_{rec} = \mathcal{L}_1 + \lambda \cdot (\mathcal{L}_{SSIM} + \mathcal{L}_{LPIPS}), \quad (4)$$

where $\lambda = 0.2$ is the weighting coefficient for structural and perceptual loss terms.

Experiments

Implementation Details

Our model is trained on a combination of datasets, including our new DRR dataset, synthetic data, and established benchmarks (*Real* and *Nature*). We employ standard data augmentation techniques and a progressive 3-stage training strategy (4 days total on 3×NVIDIA GeForce RTX 3090 GPUs) initialized with Stable Diffusion v2.1 weights. Evaluation covers three standard benchmarks plus our new DRR benchmark with 400 image pairs across 20 scenes. Our method achieves fast inference speed, processing a 768-resolution image in around 1 second on a single RTX 3090 GPU.

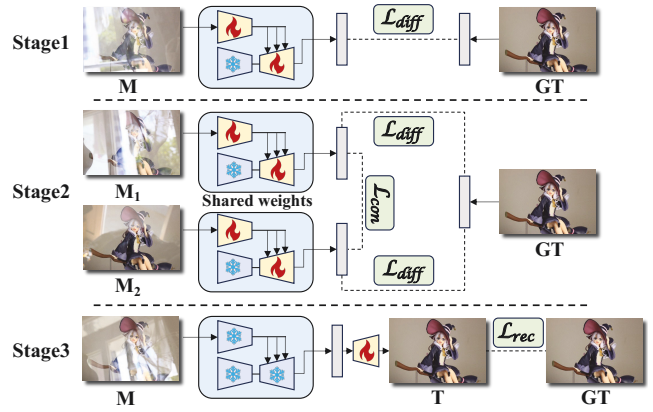


Figure 5: The three stages of progressive training. First, we train the ControlNet and the upsampling blocks of the U-Net using the basic one-step diffusion loss. Second, we finetune these components by incorporating the consistent loss. Finally, we train the cross-latent decoder using the image reconstruction loss.

Comparison to State-of-the-arts

We compare our method with several recently proposed approaches, including RobustSIRR (Song et al. 2023), DSR-Net (Hu and Guo 2023), L-Differ (Hong et al. 2024), RRW (Zhu et al. 2024), DSIT (Hu, Wang, and Guo 2024) and RDNet (Zhao et al. 2025). All evaluations are conducted using the same testing scripts and the same testing data to ensure fairness. As L-Differ’s pre-trained models and code are not available, we use the results from the original paper. Qualitative results in Fig. 6 show that our diffusion-based framework show SOTA performance. The quantitative results on the three established benchmarks are demonstrated in Table 2. Our method achieves the highest average PSNR and SSIM scores, surpassing all recent approaches. This demonstrates the overall superiority of our approach.

In addition, we evaluate the performance on our two testing benchmarks. The quantitative results are also presented in Table 2. We observe that existing methods perform poorly on these benchmarks, suggesting potential overfitting to older datasets. In contrast, our approach exhibits superior generalization capabilities, achieving strong performance on both the standard and challenging sets.

To explore the superiority of our dataset, we fine-tune all the methods in our data setting, and present a comprehensive comparison of their results before and after this adaptation in Fig. 7. It indicates that our dataset significantly enhances the robustness of reflection removal models. To illustrate this enhancement, we specifically select DSIT as an example.

Ablation Study

We conduct a comprehensive ablation study to evaluate the impact of 1) the cross-latent decoder module, 2) our new dataset DRR, and 3) the reflection-invariant finetuning strategy by adding the component gradually from a baseline. The results are summarized in Table 3 and visual quality is demonstrated in Fig. 8




Figure 6: Qualitative comparison of our method with other approaches, demonstrating superior performance in preserving background details while effectively removing reflections.



Figure 7: Visual comparison of DSIT performance before and after fine-tuning on our dataset. The results show significant improvements after adaptation to our data setting.

Cross-Latent Decoder The baseline employs a vanilla decoder from Stable Diffusion without connecting it with encoder, thus failing to restore high-frequency details. To ad-

dress the over-smooth issue of one-step diffusion process, we consider cross-latent decoder as an essential component. After incorporating this module, more original information is directly injected during the decoding process, effectively mitigating blurring effects and preserving details.

Influence of Dataset Training without our newly proposed dataset DRR, our experiments reveal that incorporating DRR enhances the performance across all benchmarks, including both established benchmarks and our proposed ones. This demonstrates that DRR introduces diverse and challenging scenarios, which are crucial for improving the generalization and robustness of reflection removal models.

Reflection-Invariant Finetuning Initially, we train the model on the entire dataset by treating all image pairs equally, without any distinction. However, experiments confirm that it is not a ideal solvment, thus leading to sub-optimal performance. After incorporating our reflection-invariant finetuning strategy, our model become more robust in handling diverse reflection scenarios, ultimately achieving more stable and reliable results.

Methods	Venue	<i>Nature</i> (20)		<i>Real</i> (20)		<i>SIR</i> ² (500)		<i>DRR-S</i> (200)		<i>DRR-C</i> (200)	
		PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
RobustSIRR	CVPR 2023	20.94	0.770	22.71	0.787	22.61	0.872	19.68	0.756	20.24	0.692
DSRNet	ICCV 2023	24.86	0.823	23.31	0.791	25.65	0.919	22.33	0.846	21.93	0.820
RRW	CVPR 2024	25.79	0.833	21.51	0.767	25.31	0.907	22.39	0.857	21.84	0.820
L-DiffER	ECCV 2024	23.95	0.831	23.77	0.821	25.18	0.911	-	-	-	-
DSIT	NeurIPS 2024	<u>26.25</u>	0.833	24.54	0.814	26.34	0.922	<u>23.48</u>	<u>0.869</u>	<u>22.46</u>	0.817
RDNet	CVPR 2025	25.94	<u>0.835</u>	<u>25.10</u>	<u>0.833</u>	<u>26.38</u>	<u>0.924</u>	23.31	0.865	22.25	<u>0.823</u>
Ours	-	27.05	0.846	25.24	0.840	27.32	0.931	27.21	0.901	23.83	0.840

Table 2: Quantitative comparisons on the existing three reflection benchmarks, and our new dataset DRR, consisting a standard set (DRR-S) and a challenging set (DRR-C). The scores after finetuning on our data setting are labeled as \dagger . The best results are in bold, and the second-best results are underlined.

CLD	DRR	RIF	<i>Nature</i> (20)		<i>Real</i> (20)		<i>SIR</i> ² (500)		<i>DRR-S</i> (200)		<i>DRR-C</i> (200)	
			PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
×	×	×	25.54	0.782	23.67	0.762	24.96	0.861	22.32	0.769	20.50	0.666
✓	×	×	26.70	0.834	23.99	0.828	25.98	0.921	24.43	0.881	22.51	0.828
✓	✓	×	26.71	0.839	24.81	0.836	26.78	0.927	26.44	0.898	23.34	0.835
✓	✓	✓	27.05	0.846	25.24	0.840	27.32	0.931	27.21	0.901	23.83	0.840

Table 3: Quantitative results of ablation study, CLD refers to Cross-Latent Decoder, DRR is our dataset Diverse Reflection Removal, RIF refers to Reflection-Invariant Finetuning.




Figure 8: Qualitative results of the ablation study. CLD is Cross-Latent Decoder, DRR is our dataset Diverse Reflection Removal, RIF refers to Reflection-Invariant Finetuning. Baseline is the basic setting without CLD, DRR and RIF.

Conclusion

We present a comprehensive framework for single-image reflection removal, comprising a high-quality dataset (DRR) featuring diverse reflection scenarios and an efficient capture pipeline, a robust diffusion-based model tailored for reflection removal, and a novel reflection-invariant finetuning strategy to enhance generalization. Our dataset significantly improves training quality by addressing the limitations of existing datasets. Our method achieves state-of-the-art performance on both common and our newly introduced bench-

marks. Extensive experiments and ablation studies validate the effectiveness of each component, demonstrating superior generalization capabilities and robustness across diverse real-world scenarios. This work not only advances reflection removal technology but also provides a valuable resource and a strong baseline for future research in the field.

Limitation When the reflection and target scene are seamlessly integrated, it is hard to distinguish which one is the target to be preserved. Incorporating semantic or user guidance offers a promising direction for future research.

Acknowledgments

This work was supported by the NSFC under Grant 62322604 and 62576207. The authors express gratitude to the anonymous reviewers for their valuable feedback.

References

- Chen, L.; Chu, X.; Zhang, X.; and Sun, J. 2022. Simple baselines for image restoration. In *European conference on computer vision*, 17–33. Springer.
- Dong, Z.; Xu, K.; Yang, Y.; Bao, H.; Xu, W.; and Lau, R. W. 2021. Location-aware single image reflection removal. In *Proceedings of the IEEE/CVF international conference on computer vision*, 5017–5026.
- Fan, Q.; Yang, J.; Hua, G.; Chen, B.; and Wipf, D. 2017. A generic deep architecture for single image reflection removal and image smoothing. In *Proceedings of the IEEE International Conference on Computer Vision*, 3238–3247.
- Guo, X.; Cao, X.; and Ma, Y. 2014. Robust separation of reflection from multiple images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2187–2194.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Hong, Y.; Zhong, H.; Weng, S.; Liang, J.; and Shi, B. 2024. L-DiffER: Single Image Reflection Removal with Language-Based Diffusion Model. In *European Conference on Computer Vision*, 58–76. Springer.
- Hu, Q.; and Guo, X. 2021. Trash or treasure? an interactive dual-stream strategy for single image reflection separation. *Advances in Neural Information Processing Systems*, 34: 24683–24694.
- Hu, Q.; and Guo, X. 2023. Single image reflection separation via component synergy. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13138–13147.
- Hu, Q.; Wang, H.; and Guo, X. 2024. Single Image Reflection Separation via Dual-Stream Interactive Transformers. *Advances in Neural Information Processing Systems*, 37: 55228–55248.
- Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35: 23593–23606.
- Kim, S.; Huo, Y.; and Yoon, S.-E. 2020. Single image reflection removal with physically-based training images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5164–5173.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4026.
- Lan, G.; Ma, Q.; Yang, Y.; Wang, Z.; Wang, D.; Li, X.; and Zhao, B. 2025. Efficient Diffusion as Low Light Enhancer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 21277–21286.
- Lei, C.; Huang, X.; Zhang, M.; Yan, Q.; Sun, W.; and Chen, Q. 2020. Polarized reflection removal with perfect alignment in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1750–1758.
- Lei, C.; Jiang, X.; and Chen, Q. 2023. Robust reflection removal with flash-only cues in the wild. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(12): 15530–15545.
- Li, C.; Yang, Y.; He, K.; Lin, S.; and Hopcroft, J. E. 2020. Single image reflection removal through cascaded refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3565–3574.
- Li, T.; Chan, Y.-H.; and Lun, D. P. 2020. Improved multiple-image-based reflection removal algorithm using deep neural networks. *IEEE Transactions on Image Processing*, 30: 68–79.
- Li, Y.; and Brown, M. S. 2014. Single image layer separation using relative smoothness. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2752–2759.
- Liu, Y.; Ma, X.; Bailey, J.; and Lu, F. 2020. Reflection backdoor: A natural backdoor attack on deep neural networks. In *Computer vision—ECCV 2020: 16th European conference, Glasgow, UK, August 23–28, 2020, proceedings, part X 16*, 182–199. Springer.
- Lowe, D. G. 2004. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60: 91–110.
- Lugmayr, A.; Danelljan, M.; Romero, A.; Yu, F.; Timofte, R.; and Van Gool, L. 2022. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11461–11471.
- Niklaus, S.; Zhang, X. C.; Barron, J. T.; Wadhwa, N.; Garg, R.; Liu, F.; and Xue, T. 2021. Learned dual-view reflection removal. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3713–3722.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PmLR.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, 234–241. Springer.
- Shih, Y.; Krishnan, D.; Durand, F.; and Freeman, W. T. 2015. Reflection removal using ghosting cues. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3193–3201.

- Song, J.; Meng, C.; and Ermon, S. 2020. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*.
- Song, Z.; Zhang, Z.; Zhang, K.; Luo, W.; Fan, Z.; Ren, W.; and Lu, J. 2023. Robust single image reflection removal against adversarial attacks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24688–24698.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1921–1930.
- Wan, R.; Shi, B.; Duan, L.-Y.; Tan, A.-H.; and Kot, A. C. 2017. Benchmarking single-image reflection removal algorithms. In *Proceedings of the IEEE International Conference on Computer Vision*, 3922–3930.
- Wan, R.; Shi, B.; Hwee, T. A.; and Kot, A. C. 2016. Depth of field guided reflection removal. In *2016 IEEE International Conference on Image Processing (ICIP)*, 21–25. IEEE.
- Wan, R.; Shi, B.; Li, H.; Duan, L.-Y.; and Kot, A. C. 2021. Face image reflection removal. *International Journal of Computer Vision*, 129: 385–399.
- Wang, J.; Yue, Z.; Zhou, S.; Chan, K. C.; and Loy, C. C. 2024a. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, 132(12): 5929–5949.
- Wang, T.; Lu, W.; Zhang, K.; Luo, W.; Kim, T.-K.; Lu, T.; Li, H.; and Yang, M.-H. 2024b. Promptrr: Diffusion models as prompt generators for single image reflection removal. *arXiv preprint arXiv:2402.02374*.
- Wang, T.; Xie, M.; Cai, H.; Shah, S.; and Metzler, C. A. 2024c. Flash-Split: 2D Reflection Removal with Flash Cues and Latent Diffusion Separation. *arXiv preprint arXiv:2501.00637*.
- Wang, Y.; Yang, W.; Chen, X.; Wang, Y.; Guo, L.; Chau, L.-P.; Liu, Z.; Qiao, Y.; Kot, A. C.; and Wen, B. 2024d. Sinsr: diffusion-based image super-resolution in a single step. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 25796–25805.
- Wang, Z.; Cun, X.; Bao, J.; Zhou, W.; Liu, J.; and Li, H. 2022. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 17683–17693.
- Wei, K.; Yang, J.; Fu, Y.; Wipf, D.; and Huang, H. 2019. Single image reflection removal exploiting misaligned training data and network enhancements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8178–8187.
- Wen, Q.; Tan, Y.; Qin, J.; Liu, W.; Han, G.; and He, S. 2019. Single image reflection removal beyond linearity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3771–3779.
- Whang, J.; Delbracio, M.; Talebi, H.; Saharia, C.; Dimakis, A. G.; and Milanfar, P. 2022. Deblurring via stochastic refinement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 16293–16303.
- Xia, B.; Zhang, Y.; Wang, S.; Wang, Y.; Wu, X.; Tian, Y.; Yang, W.; and Van Gool, L. 2023. Diffir: Efficient diffusion model for image restoration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13095–13105.
- Xie, S.; Zhang, Z.; Lin, Z.; Hinz, T.; and Zhang, K. 2023. Smartbrush: Text and shape guided object inpainting with diffusion model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 22428–22437.
- Xu, G.; Ge, Y.; Liu, M.; Fan, C.; Xie, K.; Zhao, Z.; Chen, H.; and Shen, C. 2024. What Matters When Repurposing Diffusion Models for General Dense Perception Tasks? *arXiv preprint arXiv:2403.06090*.
- Yang, J.; Gong, D.; Liu, L.; and Shi, Q. 2018. Seeing deeply and bidirectionally: A deep learning approach for single image reflection removal. In *Proceedings of the european conference on computer vision (ECCV)*, 654–669.
- Yang, L.; Kang, B.; Huang, Z.; Xu, X.; Feng, J.; and Zhao, H. 2024. Depth anything: Unleashing the power of large-scale unlabeled data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10371–10381.
- Ye, C.; Qiu, L.; Gu, X.; Zuo, Q.; Wu, Y.; Dong, Z.; Bo, L.; Xiu, Y.; and Han, X. 2024. Stablenormal: Reducing diffusion variance for stable and sharp normal. *ACM Transactions on Graphics (TOG)*, 43(6): 1–18.
- Yi, X.; Xu, H.; Zhang, H.; Tang, L.; and Ma, J. 2023. Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model. In *Proceedings of the IEEE/CVF international conference on computer vision*, 12302–12311.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3836–3847.
- Zhang, X.; Ng, R.; and Chen, Q. 2018. Single image reflection separation with perceptual losses. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4786–4794.
- Zhao, H.; Li, M.; Hu, Q.; and Guo, X. 2025. Reversible decoupling network for single image reflection removal. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 26430–26439.
- Zhong, H.; Hong, Y.; Weng, S.; Liang, J.; and Shi, B. 2024. Language-guided image reflection separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24913–24922.
- Zhou, H.; Dong, W.; Liu, X.; Zhang, Y.; Zhai, G.; and Chen, J. 2025. Low-light image enhancement via generative perceptual priors. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 10752–10760.
- Zhu, Y.; Fu, X.; Jiang, P.-T.; Zhang, H.; Sun, Q.; Chen, J.; Zha, Z.-J.; and Li, B. 2024. Revisiting SingleImage Reflection Removal in the Wild. In *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.