LENSBRIDGE: FREQUENCY-GUIDED COMPOUND DEGRADATION ADAPTATION FOR LENS ABERRATION CORRECTION AND VEILING GLARE REMOVAL

Anonymous authors
Paper under double-blind review

ABSTRACT

Simplified optical systems often exhibit residual lens aberrations and Veiling Glare (VG), resulting in spatially varying blur and contrast reduction. Large-scale Lens Libraries (LensLib) enable reusable aberration correction models by covering diverse Point Spread Functions (PSFs), but their aberration-only training distribution does not include target-specific veiling glare. Extending such foundations to compound degradation is challenging because realistic target-system compound pairs are difficult to obtain. To address this challenge, we propose LENSBRIDGE, a two-stage framework that first establishes a reusable aberration correction foundation and then adapts it to compound optical degradation using only a few unpaired target observations. In Stage I, we build a PSF-aware one-step diffusion foundation by constructing discrete degradation priors from LensLib PSFs and learning to retrieve them directly from aberrated images, enabling PSF-aware correction without requiring explicit PSF at inference. In Stage II, we adapt this foundation to compound degradation through frequency-domain guidance. Our analysis suggests that aberrations affect a broad frequency range, whereas additional veiling glare exhibits stronger low-frequency amplitude variations, motivating a shared target-domain Fourier-amplitude statistic. At the data level, Frequency-guided Degradation Completion (FDC) transfers target low-frequency characteristics to LensLib aberrated images while preserving aberration structures to synthesize compound training pairs; at the model level, Frequency-guided Pseudo Decomposition (FPD) forms aberration- and VG-dominant pseudo observations to condition separate adaptation branches. Extensive experiments across multiple optical systems demonstrate that LENSBRIDGE effectively extends reusable aberration correction foundations to joint aberration correction and veiling glare removal without target-system paired supervision.

1 INTRODUCTION

Simplified optical designs enable compact imaging systems for mobile photography (Chen et al., 2021a; 2023) and emerging meta-optical devices (Li et al., 2021b; Chen et al., 2024), but often leave substantial residual lens aberrations. Meanwhile, surface imperfections and non-ideal coatings can also induce diffuse stray-light scattering, manifested as *Veiling Glare* (VG) (Qian et al., 2026b). Lens aberrations introduce spatially varying blur and chromatic dispersion, while veiling glare causes diffuse luminance contamination and contrast reduction. Their coexistence results in *compound optical degradation*, motivating computational restoration capable of addressing both effects.

Existing computational restoration for optical degradation has primarily focused on lens aberrations. Conventional learning-based approaches typically rely on target-system synthetic pairs generated from Point Spread Functions (PSFs) (Chen et al., 2021a;b). As illustrated by the *synthetic-to-real* paradigm in Fig. 1(a), such models are trained on target-lens synthetic data and then transferred to real aberrated images (Jiang et al., 2025). Acquiring target-lens PSFs and constructing synthetic training pairs for new optical system make deployment costly. Pretraining on large-scale Lens Libraries (LensLib) alleviates this burden by learning reusable correction models across diverse aberration distributions (Gong et al., 2024; Jiang et al., 2026; Qian et al., 2026a). Among these methods,

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

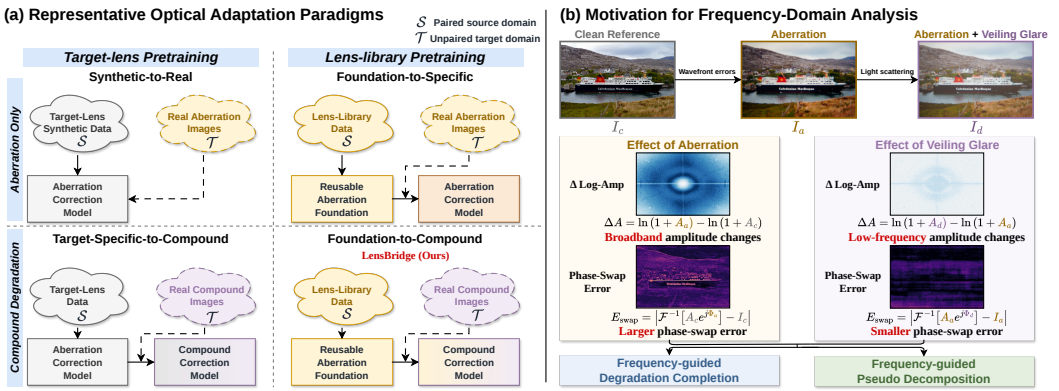

Figure 1: **Motivation for LensBridge.** (a) Comparison of optical adaptation paradigms in terms of source pretraining and target degradation. LensBridge adapts a reusable aberration correction foundation to target compound degradation using only unpaired target images. (b) In the illustrated example, aberration induces broadband amplitude changes, whereas the additional veiling glare produces predominantly low-frequency amplitude changes and a smaller phase-swap reconstruction error in the illustrated example, motivating the frequency-guided designs of FDC and FPD.

OmniLens (Jiang et al., 2026) follows the *foundation-to-specific* paradigm in Fig. 1(a), adapting a LensLib foundation to unseen lenses with unpaired aberrated images and a dark-channel prior.

However, these reusable foundations are learned within the aberration-only degradation space of LensLib and therefore cannot account for veiling glare. Extending them to compound degradation remains challenging, as physically accurate veiling glare simulation requires optomechanical models and computationally expensive non-sequential ray tracing (Raskar et al., 2008; Hullin et al., 2011), making reliable target-system compound pairs difficult to obtain. A recent solution is target-specific compound degradation synthesis. DeVeiler (Qian et al., 2026b) follows the *target-specific-to-compound* paradigm in Fig. 1(a), using target-lens aberration pairs and unpaired compound observations to synthesize realistic lens-specific compound degradations for joint restoration. However, its synthesis remains tailored to each target system rather than reusing a aberration correction foundation across lenses. In contrast, we pursue a *foundation-to-compound* paradigm, where reusable aberration knowledge is retained from LensLib and only a few unpaired compound observations are required to adapt the foundation to a new optical system.

To enable such adaptation, we examine the two degradation mechanisms in the frequency domain, where PSF-induced blur and diffuse stray-light effects exhibit different spectral responses. PSF-induced aberration modifies image content across a broad frequency range, whereas veiling glare primarily affects low-frequency amplitude components, as illustrated in Fig. 1(b). This distinction motivates using target-domain low-frequency amplitude statistics to characterize the degradation information missing from the aberration-only foundation.

We therefore propose LensBridge, a two-stage framework that first establishes a reusable aberration correction foundation and then adapts it to compound optical degradation. In Stage I, we establish a PSF-aware one-step diffusion foundation for reusable aberration correction. LensLib PSFs are encoded into a discrete degradation prior space, while image-conditioned retrieval enables the corresponding priors to be inferred directly from aberrated observations. The retrieved priors condition the diffusion restoration backbone, retaining PSF-aware degradation guidance without requiring explicit PSFs at inference.

In Stage II, we adapt this foundation to target compound degradation using only a few unpaired target observations. Based on the above frequency distinction, we estimate a shared target-system Fourier-amplitude statistic to guide adaptation at both the data and model levels. At the data level, Frequency-guided Degradation Completion (FDC) performs data-level adaptation by transferring target-domain low-frequency degradation characteristics to LensLib aberrated images while preserving their aberration-related structures and clean counterparts, thereby constructing target-like compound training pairs without learning an additional degradation generator. At the model level, Frequency-guided Pseudo Decomposition (FPD) uses the same frequency statistic to construct aberration-dominant and VG-dominant pseudo observations, allowing the two adaptation branches to model their distinct degradation characteristics separately. Together, FDC provides target-like

training supervision, while FPD separately conditions the aberration and VG adaptation branches. Our main contributions are summarized as follows:

- We propose LENSBRIDGE, a two-stage framework that first establishes a reusable LensLib-based aberration correction foundation and then adapts it to compound aberration and veiling-glare degradation using only a few unpaired target observations.
- We introduce FDC and FPD, which leverage a shared target-domain Fourier-amplitude statistic to construct target-like compound training pairs and aberration-/VG-dominant pseudo observations, respectively.
- Extensive experiments across multiple optical systems demonstrate the effectiveness of LENSBRIDGE for joint aberration correction and veiling glare removal.

2 RELATED WORK

Lens Aberration Correction. Computational Aberration Correction (CAC) aims to restore images degraded by residual lens aberrations. Early approaches relied on deconvolution or PSF-based inverse filtering (Wiener, 1949; Fish et al., 1995; Schuler et al., 2012). With the emergence of deep learning, learning-based CAC methods (Chen et al., 2021a;b; Jiang et al., 2024; Luo et al., 2024a; Zhang et al., 2025; Mandal et al., 2025) learn restoration mappings from simulated optical degradations and achieve improved performance. However, these methods typically depend on lens-specific simulation pipelines and suffer from limited generalization to unseen optical systems. Recent works explore reusable CAC models through large-scale lens library pre-training (Jiang et al., 2026; Qian et al., 2026a), where diverse aberration distributions improve cross-lens generalization and support adaptation to new lenses. Nevertheless, existing LensLib-based frameworks remain confined to PSF-induced aberrations and do not address compound degradations involving veiling glare.

Veiling Glare Removal. Veiling glare is a diffuse degradation caused by stray-light scattering inside optical systems, leading to luminance contamination and contrast reduction. Unlike lens flare, which typically appears as localized structures such as bright spots and streaks (Wu et al., 2021; Dai et al., 2022; 2023; 2024; Vasluianu et al., 2024; Zhu & Lee, 2025), veiling glare is spatially distributed and dominated by low-frequency variations. Traditional approaches rely on physical modeling or image decomposition, including Glare Spread Function (GSF) based inverse filtering (Talvala et al., 2007) and glare-component decomposition (Zhang et al., 2018b), but often require restrictive assumptions, calibration, or multiple observations. Recent data-driven methods (Shoshin & Shvets, 2021) use simplified flare or scattering models for training-data synthesis, leaving a substantial gap from lens-specific real degradations. DeVeiler (Qian et al., 2026b) addresses this limitation through target-specific compound degradation synthesis, but its degradation modeling must be repeated for each new optical system. In contrast, LENSBRIDGE reuses a LensLib-based aberration correction foundation and adapts it to compound degradation using only unpaired target observations.

Frequency-Guided Degradation Adaptation. Frequency-domain statistics have been widely explored for domain adaptation and image restoration. FDA transfers low-frequency Fourier amplitudes between source and target domains for unsupervised semantic segmentation (Yang & Soatto, 2020), while FCA exploits frequency consistency to reduce the synthetic-to-real gap in real-world super-resolution (Ji et al., 2021). Frequency-aware dehazing methods further exploit amplitude-spectrum discrepancies between synthetic and real degradations for adaptation or restoration (Yu et al., 2022; Liu et al., 2025). These methods demonstrate the effectiveness of frequency statistics for characterizing domain shifts, but mainly use them for global source-to-target alignment or adaptation within a single degradation type. In contrast, LENSBRIDGE uses an aggregated target-domain Fourier-amplitude statistic to complete the missing target degradation characteristics associated with VG while preserving the existing LensLib aberration structure, and further reuses the same statistic to construct aberration-dominant and VG-dominant pseudo observations.

3 METHODOLOGY

3.1 PROBLEM FORMULATION

We aim to recover a latent clean image I_c from a compound-degraded image I_{de} captured by a target optical system. Following (Qian et al., 2026b), we model the joint effects of spatially varying lens

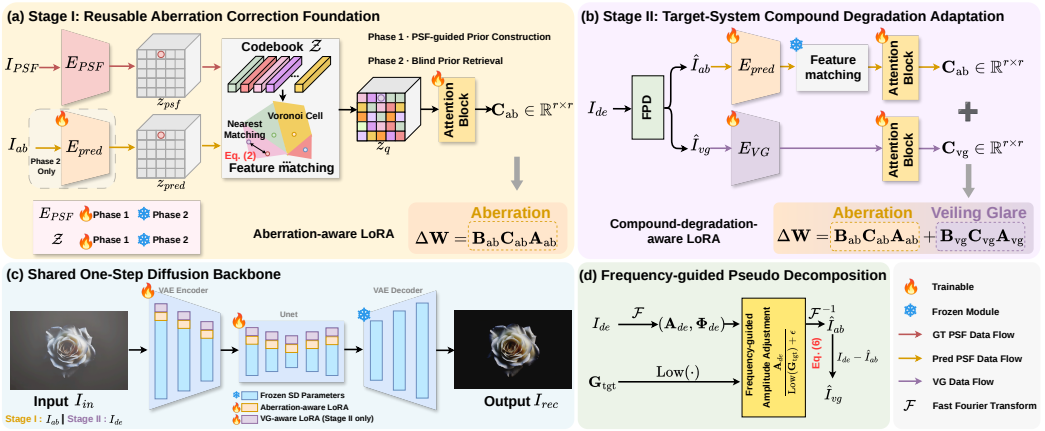

Figure 2: **Overview of LensBridge.** LensBridge consists of two stages: (a) a reusable PSF-aware aberration correction foundation, where discrete degradation priors are constructed from LensLib PSFs and retrieved directly from aberrated observations; and (b) target-system compound degradation adaptation, where aberration-dominant and VG-dominant pseudo observations separately condition the two adaptation branches. The shared one-step diffusion backbone and the frequency-guided pseudo decomposition mechanism are detailed in (c) and (d), respectively.

aberration and veiling glare. Approximating the PSF as shift-invariant within each spatial patch (Li et al., 2021a; Luo et al., 2024a; Yang et al., 2025), we express the degradation independently for each color channel as:

$$I^p_{de} = \underbrace{(I^p_c \otimes K^p)}_{I^p_{ab}} \odot T^p + I^p_g, \tag{1}$$

where p indexes spatial patches, $\otimes$ denotes convolution, and $\odot$ denotes element-wise multiplication. K^p is the local PSF, I^p_{ab} is the aberrated patch, T^p models contrast attenuation, and I^p_g represents additive glare (Zhang et al., 2018b; Shoshin & Shvets, 2021). The color-channel index is omitted for clarity. We use a LensLib source dataset $\mathcal{D}_s = \{(I^s_{ab,i}, I^s_{c,i})\}_{i=1}^{N_s}$, with a corresponding PSF for each pair, and a small set of unpaired compound-degraded target images $\mathcal{D}_t = \{I^t_{de,j}\}_{j=1}^{N_t}$. Here, s and t denote the source and target domains, respectively, and $N_t \ll N_s$. Our objective is to first establish a reusable aberration correction foundation from $\mathcal{D}_s$ and then adapt it to target compound degradation using only the unpaired observations in $\mathcal{D}_t$.

3.2 Framework Overview

As illustrated in Fig. 2, LensBridge consists of two stages: establishing a reusable aberration correction foundation and adapting it to target-system compound degradation. We adopt SD-Turbo (Sauer et al., 2024), distilled from Stable Diffusion (Rombach et al., 2022), as the one-step diffusion backbone. In Stage I, we construct discrete aberration priors from LensLib PSFs and train an image-based predictor to retrieve them from aberrated images. These priors condition low-rank adaptation (LoRA) (Hu et al., 2022) of the diffusion backbone. In Stage II, we estimate a shared target-system Fourier-amplitude statistic from a few unpaired target observations. Frequency-guided Degradation Completion (FDC) uses this statistic to construct target-like compound training pairs, while Frequency-guided Pseudo Decomposition (FPD) constructs aberration-dominant and VG-dominant pseudo observations to condition separate adaptation branches.

3.3 Stage I: Reusable Aberration Correction Foundation

Aberration correction benefits from explicit PSF guidance, whereas target-system PSFs are typically unavailable at inference. We therefore construct a PSF-derived discrete prior space and train the Stage-I foundation in two phases. In Phase 1, *PSF-guided Prior Construction*, source-domain PSFs are encoded into discrete aberration priors to condition the diffusion backbone. In Phase 2, *Blind Prior Retrieval*, this prior space is fixed and an image-based predictor learns to retrieve the corresponding priors from aberrated observations by matching its predictions to PSF-derived representations during training, enabling PSF-aware correction without requiring PSFs at inference.

Phase 1: PSF-guided Prior Construction. Following the discrete PSF representation strategy of FoundCAC (Qian et al., 2026a), we employ vector quantization to construct the PSF-derived prior space. Specifically, a PSF encoder E_{PSF} maps a PSF representation $I_{\mathrm{PSF}} \in \mathbb{R}^{H\times W\times N_p}$ to a continuous representation $\mathbf{z}_{\mathrm{psf}} \in \mathbb{R}^{H_z\times W_z\times C_z}$, which is quantized using a learnable codebook $\mathcal{Z}$ to obtain $\mathbf{z}^q_{\mathrm{psf}} \in \mathbb{R}^{H_z\times W_z\times C_z}$. The resulting codebook defines a discrete PSF-derived prior space for PSF-guided conditioning in Phase 1 and image-based retrieval in Phase 2, where image-predicted representations are matched to the corresponding PSF-derived priors during training. We regularize the encoder and codebook with the vector-quantization objective:

$$\mathcal{L}_{\mathrm{VQ}} = \left\|\mathrm{sg}[\mathbf{z}_{\mathrm{psf}}] - \mathbf{z}^q_{\mathrm{psf}}\right\|_2^2 + \beta\left\|\mathbf{z}_{\mathrm{psf}} - \mathrm{sg}[\mathbf{z}^q_{\mathrm{psf}}]\right\|_2^2, \tag{2}$$

where $\mathrm{sg}[\cdot]$ denotes stop-gradient and β controls the commitment term. The first term updates the selected codebook representations toward the encoder outputs, while the second encourages the encoder outputs to remain close to their assigned codes.

Building on degradation-aware LoRA (Zhang et al., 2024), we condition the low-rank update on PSF-derived discrete priors. For a pretrained weight $\mathbf{W}_0 \in \mathbb{R}^{d_{\mathrm{out}}\times d_{\mathrm{in}}}$, standard LoRA uses two low-rank matrices $\mathbf{B}_{\mathrm{ab}} \in \mathbb{R}^{d_{\mathrm{out}}\times r}$ and $\mathbf{A}_{\mathrm{ab}} \in \mathbb{R}^{r\times d_{\mathrm{in}}}$, where $r \ll \min(d_{\mathrm{out}}, d_{\mathrm{in}})$. An attention block maps the quantized prior $\mathbf{z}^q_{\mathrm{psf}}$ to a PSF-conditioned modulation matrix $\mathbf{C}_{\mathrm{ab}} \in \mathbb{R}^{r\times r}$, yielding:

$$\mathbf{W} = \mathbf{W}_0 + \mathbf{B}_{\mathrm{ab}}\mathbf{C}_{\mathrm{ab}}\mathbf{A}_{\mathrm{ab}}. \tag{3}$$

Compared with the standard LoRA update $\mathbf{B}_{\mathrm{ab}}\mathbf{A}_{\mathrm{ab}}$, $\mathbf{C}_{\mathrm{ab}}$ dynamically modulates the low-rank update according to the supplied aberration prior, enabling PSF-conditioned adaptation. Image reconstruction is supervised by a weighted combination of L1, perceptual, and adversarial losses, denoted as $\mathcal{L}_{\mathrm{rec}}$. Jointly optimizing $\mathcal{L}_{\mathrm{rec}}$ and $\mathcal{L}_{\mathrm{VQ}}$ encourages the discrete PSF-derived priors to provide effective conditioning for aberration restoration while regularizing their codebook assignments.

Phase 2: Blind Prior Retrieval. To enable image-only inference while retaining PSF-derived degradation guidance, we train an RGB degradation predictor E_{pred} to retrieve the discrete priors established in Phase 1 directly from an aberrated image I_{ab}. To supervise this retrieval, the PSF encoder E_{PSF} and codebook $\mathcal{Z}$ are kept frozen. During training, the corresponding PSF representation is encoded and quantized to obtain the teacher prior $\mathbf{z}^q_{\mathrm{psf}}$. A feature-matching loss $\mathcal{L}_{\mathrm{FM}}$ (Chen et al., 2022) aligns the predictor's pre-quantization feature $\mathbf{z}_{\mathrm{pred}}$ with $\mathbf{z}^q_{\mathrm{psf}}$. The predicted feature is then quantized by the frozen codebook to obtain $\mathbf{z}^q_{\mathrm{pred}}$ for restoration conditioning. The retrieval stage is optimized with $\mathcal{L}_{\mathrm{rec}} + \mathcal{L}_{\mathrm{FM}}$. At inference, E_{pred} retrieves the aberration prior directly from the observed image without requiring the PSF.

3.4 Stage II: Target-System Compound Degradation Adaptation

In Stage I, we learn a reusable aberration correction foundation for PSF-induced degradation from LensLib. However, the source distribution does not cover target-specific factor (*i.e.*, veiling glare) beyond the PSF space. The two degradations arise from different physical mechanisms and therefore exhibit distinct spectral responses. As illustrated in Fig. 1(b), PSF-induced aberration modulates a broad frequency range, whereas the spatially smooth contamination of veiling glare mainly affects low-frequency amplitudes. We therefore estimate a shared target-system Fourier-amplitude statistic from a few unpaired observations to guide data-level degradation completion and model-level pseudo decomposition.

Frequency-guided Degradation Completion. Directly transferring the Fourier amplitude of a single target image is unreliable because its spectrum is entangled with scene content. To obtain a target-system cue less sensitive to image-specific content, we aggregate Fourier-amplitude statistics over a few unpaired target observations. The resulting statistic is then used to complete the missing target degradation on LensLib aberrated images rather than regenerate the full degradation from clean images. Given a small target set $\mathcal{D}_t$ and an auxiliary clear-image collection $\mathcal{D}_{\mathrm{clr}}$, we first compute their channel-wise mean Fourier amplitudes:

$$\bar{\mathbf{A}}^c_{\mathcal{D}} = \frac{1}{N_{\mathcal{D}}}\sum_{i=1}^{N_{\mathcal{D}}}\left|\mathcal{F}(I^c_i)\right|, \qquad \mathbf{G}^c_{\mathrm{tgt}} = \frac{\bar{\mathbf{A}}^c_{\mathcal{D}_t}}{\bar{\mathbf{A}}^c_{\mathcal{D}_{\mathrm{clr}}} + \epsilon}, \tag{4}$$

where $\mathcal{F}(\cdot)$ denotes the centered 2D Fourier transform, c indexes the color channel, and ϵ ensures numerical stability. Averaging across images reduces image-specific spectral variation, allowing $\mathbf{G}_{\text{tgt}}$ to capture the Fourier-amplitude discrepancy shared by the target observations. LensLib aberrated images I^s_{ab} already contain diverse aberration characteristics that should be retained during target adaptation. We therefore modulate the Fourier amplitude of I^s_{ab} using $\mathbf{G}_{\text{tgt}}$ and retain its original phase. The low-frequency component of the modulated reconstruction is then combined with the original high-frequency component through wavelet decomposition (Wang et al., 2024):

$$\tilde{I}^s_{de} = \mathcal{W}_H(I^s_{ab}) + \mathcal{W}_L\left(\mathcal{F}^{-1}\left[(\mathbf{A}^s_{ab} \odot \mathbf{G}_{\text{tgt}})\, e^{j\mathbf{\Phi}^s_{ab}}\right]\right), \tag{5}$$

where $\mathbf{A}^s_{ab}$ and $\mathbf{\Phi}^s_{ab}$ denote the Fourier amplitude and phase of I^s_{ab}, respectively. During training, we further apply random spatial masks to modulate the transferred target characteristics, increasing spatial degradation diversity. Keeping the clean reference I^s_c unchanged yields the target-like training pair $(\tilde{I}^s_{de}, I^s_c)$, completing the missing target degradation while preserving the paired supervision from LensLib.

Frequency-guided Pseudo Decomposition. Since aberration and veiling glare exhibit distinct degradation characteristics, a shared LoRA condition may entangle their adaptation signals. We therefore construct aberration-dominant and VG-dominant pseudo observations to condition separate LoRA branches. For a compound-degraded image I_{de}, we suppress the target-specific low-frequency amplitude shift while retaining its original phase:

$$\hat{I}_{ab} = \mathcal{F}^{-1}\left[\frac{\mathbf{A}_{de}}{\text{Low}(\mathbf{G}_{\text{tgt}}) + \epsilon} e^{j\mathbf{\Phi}_{de}}\right], \qquad \hat{I}_{vg} = I_{de} - \hat{I}_{ab}, \tag{6}$$

where $\text{Low}(\cdot)$ preserves the target modulation within the low-frequency region and sets the remaining frequencies to one. The resulting $\hat{I}_{ab}$ and $\hat{I}_{vg}$ serve as aberration-dominant and VG-dominant pseudo observations rather than physically exact degradation components, and are only used to provide factor-specific adaptation cues. The aberration-dominant observation $\hat{I}_{ab}$ is processed by the Stage-I degradation predictor, and the retrieved discrete aberration prior is mapped by the attention block to $\mathbf{C}_{\text{ab}}$. In parallel, a lightweight VG encoder E_{VG} maps $\hat{I}_{vg}$ to the VG condition $\mathbf{C}_{\text{vg}} \in \mathbb{R}^{r\times r}$. To preserve these distinct adaptation signals while reusing the Stage-I aberration knowledge, we initialize the aberration branch from Stage I and introduce a separate VG branch:

$$\mathbf{W} = \mathbf{W}_0 + \mathbf{B}_{\text{ab}}\mathbf{C}_{\text{ab}}\mathbf{A}_{\text{ab}} + \mathbf{B}_{\text{vg}}\mathbf{C}_{\text{vg}}\mathbf{A}_{\text{vg}}, \tag{7}$$

where $(\mathbf{A}_{\text{vg}}, \mathbf{B}_{\text{vg}})$ are the learnable low-rank matrices of the VG branch. During Stage II, FDC-generated compound pairs are jointly used with the original LensLib aberration-only pairs for adaptation. The pretrained diffusion backbone and discrete codebook remain frozen during Stage II.

4 Experiments

4.1 Experiment Settings

Datasets. We train the aberration correction foundation on LensLib (Sun et al., 2026), which provides paired aberration data and PSFs. For target adaptation and evaluation, we use the SL and MRL systems from DeVeiler (Qian et al., 2026b), representing a large-aperture single lens and a metasurface-refractive hybrid lens, respectively. The `Screen` domain provides controlled monitor captures with 50 adaptation images and $42/25$ paired test images for SL/MRL, while the `Real` domain contains in-the-wild captures with 50 adaptation images and $51/11$ unpaired test images. All images are 1280×1920. We further evaluate on the commercial minimalist lens CAYE $50\,\text{mm}$ $f/1.4$ and a nano-optics imaging system (Tseng et al., 2021). Our default *unpaired-target setting* uses only unpaired target images. We also consider a *target aberration-paired setting* with $170/125$ additional aberration-only pairs for SL/MRL to compare with methods requiring target-specific supervision. For the generic clear-image collection $\mathcal{D}_{\text{clr}}$ in Eq. 4, we use the DIV2K training set.

Implementation Details. We use SD-Turbo (Sauer et al., 2024) as the one-step restoration backbone, with rank-16 aberration/VG LoRA branches in the VAE encoder and rank-32 branches in the diffusion UNet. In Stage I, each phase is trained for 30K iterations at 2×10^{-5}; in Stage II, we adapt to each target system for 3K iterations at 1×10^{-5}. Training uses 512×512 random crops with batch size 8 on an NVIDIA A800 GPU. Further details are provided in the supplementary material.

Table 1: **Quantitative comparisons on the SL and MRL optical systems under the unpaired target setting.** Methods requiring adaptation use only unpaired target observations. **Red** and blue indicate the best and the second best performance, respectively. * denotes the pretrained model.

Dataset	Metric	Fast Two-Step*	DA-CLIP*	AdaIR*	FoundIR-v2*	FAPE-IR*	OmniLens	DeVeiler	LensBridge (Ours)
Screen SL	PSNR ↑	16.62	16.74	17.24	17.40	17.72	17.67	19.18	**20.46**
	SSIM ↑	0.539	0.620	0.597	0.539	0.626	0.541	**0.678**	0.652
	LPIPS ↓	0.676	0.428	0.582	0.465	0.487	0.430	0.327	**0.266**
	CLIP-IQA ↑	0.298	0.389	0.246	0.476	0.253	0.609	0.526	**0.620**
	NIQE ↓	6.908	4.630	8.661	5.373	6.508	4.338	4.856	**3.417**
	MUSIQ ↑	21.86	48.48	22.54	52.45	32.73	55.66	57.96	**65.30**
	Q-Align ↑	2.567	3.122	2.366	3.272	2.932	3.859	4.154	**4.473**
Screen MRL	PSNR ↑	14.94	14.94	18.43	17.74	17.23	16.07	15.46	**19.72**
	SSIM ↑	0.610	0.641	**0.705**	0.634	0.681	0.666	0.663	0.652
	LPIPS ↓	0.503	0.484	0.466	0.479	0.472	0.453	0.449	**0.333**
	CLIP-IQA ↑	0.404	0.309	0.284	0.351	0.321	0.361	0.456	**0.556**
	NIQE ↓	4.310	4.360	4.582	5.062	5.208	4.376	4.908	**3.349**
	MUSIQ ↑	44.49	42.64	42.22	43.23	42.15	45.75	48.22	**60.55**
	Q-Align ↑	3.384	3.258	3.494	3.423	3.423	3.394	3.582	**4.400**
Real SL	CLIP-IQA ↑	0.356	0.370	0.274	0.370	0.310	0.464	0.466	**0.572**
	NIQE ↓	5.817	5.465	7.146	6.231	6.155	4.794	5.201	**4.150**
	MUSIQ ↑	33.58	38.07	25.86	40.14	36.57	51.73	54.33	**60.32**
	Q-Align ↑	2.942	3.261	2.625	3.103	3.321	3.946	4.290	**4.354**
Real MRL	CLIP-IQA ↑	0.388	0.297	0.342	0.347	0.314	0.350	0.422	**0.536**
	NIQE ↓	5.332	6.658	5.802	6.169	5.947	5.299	5.629	**4.363**
	MUSIQ ↑	33.35	38.44	34.27	36.28	36.52	38.91	43.79	**56.30**
	Q-Align ↑	2.687	2.612	2.762	2.800	2.932	2.720	2.846	**3.919**

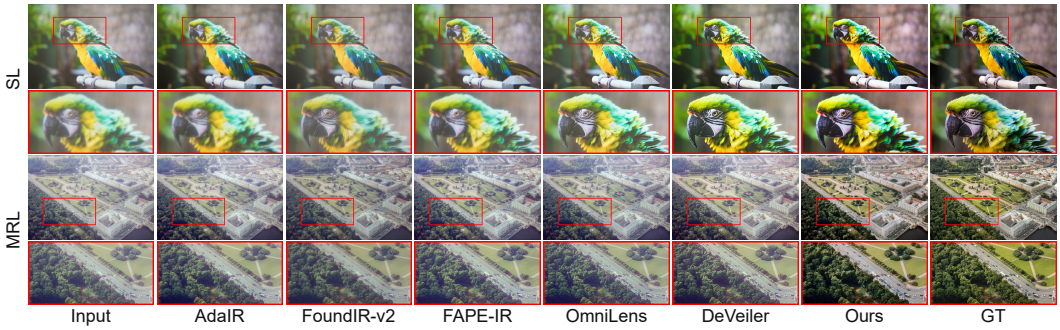

Figure 3: **Visual results on `Screen` domain.** Zoom in for the best view.

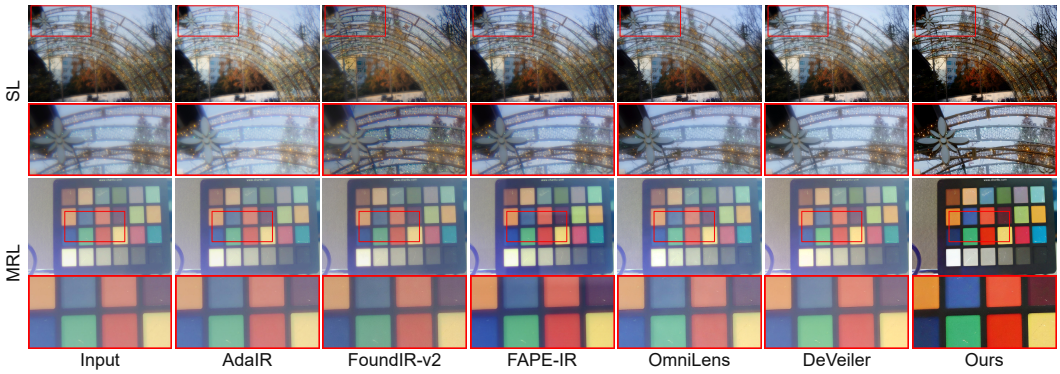

Figure 4: **Visual results on `Real` domain.** Zoom in for the best view.

Evaluation Metrics. We report PSNR, SSIM, and LPIPS (Zhang et al., 2018a) on the paired `Screen` domain . For both domains, we report CLIP-IQA (Wang et al., 2023), NIQE (Mittal et al., 2013), MUSIQ (Ke et al., 2021), and Q-Align (Wu et al., 2024a) as no-reference quality metrics.

4.2 Comparison with State-of-the-Art Methods

We evaluate LensBridge under two target-adaptation settings. Our primary setting uses only unpaired compound-degraded target images and is also evaluated on additional optical systems. For

Table 2: **Quantitative comparisons on the SL and MRL optical systems in the `Screen` domain under the target aberration-paired setting.** **Red** and blue indicate the best and the second-best performance, respectively. * denotes using the pretrained model.

Dataset	Metric	SwinIR	SwinIR+ Flare7K++*	SwinIR+ DiffDehaze*	DiffDehaze	QDMR	OmniLens	DeVeiler	LensBridge (Ours)
`Screen` SL	PSNR ↑	18.18	21.67	19.31	18.87	18.45	20.53	22.38	**22.69**
	LPIPS ↓	0.298	0.297	0.347	0.435	0.291	0.262	0.261	**0.202**
	MUSIQ ↑	62.08	59.60	64.29	39.35	61.52	65.21	62.29	**68.20**
	Q-Align ↑	4.091	4.360	4.318	3.843	4.087	4.221	4.299	**4.709**
`Screen` MRL	PSNR ↑	19.34	20.74	18.28	18.55	20.67	20.64	**21.57**	21.46
	LPIPS ↓	0.354	0.336	0.392	0.426	0.315	0.300	0.301	**0.266**
	MUSIQ ↑	58.72	58.22	64.70	54.96	61.29	62.84	58.43	**67.74**
	Q-Align ↑	4.206	4.356	4.375	4.343	4.344	4.435	4.310	**4.617**

fair comparison with DeVeiler (Qian et al., 2026b) and its benchmarked methods, we additionally consider the target aberration-paired setting.

Unpaired target setting. In the primary setting, target adaptation uses only unpaired compound-degraded images. We compare against three groups of baselines. The first is Fast Two-Step (Eboli et al., 2022), a blind aberration correction method. The second comprises four general image restoration methods: DA-CLIP (Luo et al., 2024b), AdaIR (Cui et al., 2025), FoundIR-v2 (Chen et al., 2026), and FAPE-IR (Liu et al., 2026). Their pretrained models are directly used for inference. The third comprises two domain adaptation methods: OmniLens (Jiang et al., 2026) for aberration correction and DeVeiler (Qian et al., 2026b) for joint aberration and VG removal. Both are evaluated under the same unpaired-target setting as LensBridge. As shown in Tab. 1, LensBridge performs favorably against existing methods on both `Screen` and `Real` domains. The gains are particularly consistent in LPIPS and no-reference image quality metrics, including over DeVeiler under the same setting. Visual comparisons in Figs. 3 and 4 further show that LensBridge restores fine details from aberration blur and color contrast degraded by veiling glare on both domains. These results demonstrate effective adaptation of a reusable aberration foundation to compound degradation with unpaired target observations.

Beyond SL and MRL, we further evaluate LensBridge on the CAYE lens and nano-optics system, independently adapting the foundation to each system with only a few unpaired degraded observations. As shown in Fig. 5, Stage I already recovers details degraded by aberration, while target adaptation further suppresses veiling glare and restores color and contrast. These results demonstrate the applicability of the two-stage framework across diverse optical systems.

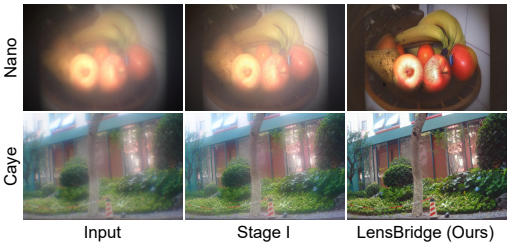

Figure 5: **Results on additional optical systems.** LensBridge effectively adapts to the CAYE lens and nano-optics system using only a few unpaired target observations.

Target aberration-paired setting. For comparison with methods benchmarked in DeVeiler (Qian et al., 2026b), we additionally evaluate the target aberration-paired setting, where aberration-clean pairs from the target optical system are available. We compare with cascaded pipelines combining a retrained SwinIR (Liang et al., 2021) with Flare7K++ (Dai et al., 2024) or DiffDehaze (Wang et al., 2025), as well as DiffDehaze, QDMR (Jiang et al., 2025), OmniLens (Jiang et al., 2026), and DeVeiler (Qian et al., 2026b). As shown in Tab. 2, LensBridge achieves the best overall performance across both optical systems and domains. The cascaded pipelines remain limited, while FDC bridges the target degradation gap and FPD separately conditions the aberration and VG LoRA branches. Together with the unpaired-target results, these results demonstrate the effectiveness of LensBridge under different levels of target-lens supervision. Additional visual results are provided in the appendix.

4.3 Ablation Studies

Effectiveness of FDC. We adopt the regression restoration model from DeVeiler (Qian et al., 2026b) and pretrain it on aberration-only pairs from the `Screen`-SL dataset. We compare different synthesis strategies, including CycleGAN (Zhu et al., 2017), HazeGen (Wang et al., 2025), VeilGen (Qian et al., 2026b), and FDA (Yang & Soatto, 2020), under the same adaptation setting. As shown in Tab. 3, FDC achieves the best performance by using aggregated target-domain spectral statis-

Method	PSNR↑	LPIPS↓	MUSIQ↑	Q-Align↑
CycleGAN	18.85	0.319	56.46	3.998
HazeGen	20.82	0.269	**63.26**	4.229
VeilGen	21.56	0.263	62.87	4.244
FDA-style	20.95	0.352	55.35	3.955
Ours	**23.82**	**0.253**	61.33	**4.336**

Table 3: Comparison of data synthesis strategies on the `Screen` SL domain.

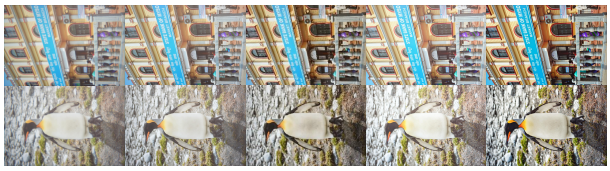

Figure 6: Generality of FDC across restoration backbones on the `Screen` SL domain.

Exp.	Steps	FDC	AB LoRA	VG LoRA	FPD	PSNR↑	LPIPS↓	MUSIQ↑	Q-Align↑
(a)	30K					15.075	0.332	53.651	3.725
(b)	30K	✓				19.722	0.314	54.937	4.055
(c)	60K	✓				19.386	0.299	58.057	4.151
(d)	60K	✓	✓			20.444	0.268	64.064	4.400
(e)	60K	✓	✓	✓		20.316	0.269	65.286	4.458
(f)	60K	✓	✓	✓	✓	**20.462**	**0.266**	**65.303**	**4.473**

Table 4: Component analysis of LENSBRIDGE.

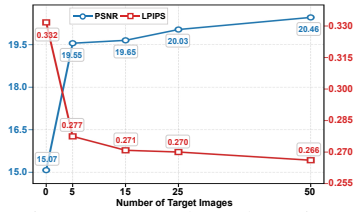

Figure 7: Few-shot adaptation.



Figure 8: Analysis of target spectral statistics.

tics rather than directly transferring a single target observation. We further evaluate FDC with SwinIR (Liang et al., 2021) and S3Diff (Zhang et al., 2024). As shown in Fig. 6, FDC consistently improves both models, demonstrating its generality across restoration architectures.

Effectiveness of Individual Components. We conduct component ablations on the `Screen`-SL dataset. As shown in Tab. 4, FDC-based adaptation in (b) substantially improves the baseline, while longer backbone training in (c) brings limited gains. The aberration LoRA branch in (d) further improves restoration quality, whereas adding VG LoRA without FPD in (e) provides limited improvement under entangled degradations, suggesting that additional adaptation capacity alone cannot resolve the coupling between different degradation factors. With FPD, (f) enables separate degradation guidance for the two LoRA branches and achieves the best overall performance.

Few-shot Adaptation. We evaluate LENSBRIDGE's data efficiency on the `Screen` SL dataset by varying the number of unpaired target images. In Fig. 7, a few target observations yield substantial gains over the non-adapted model, while additional images yield diminishing improvements.

Analysis of Target Spectral Statistics. We analyze the $\mathbf{G}_{\text{tgt}}$ on the `Real` MRL dataset. In Fig. 8(a), aggregating more target observations stabilizes the spectral statistic. Using this statistic, FDC in Fig. 8(b) completes target degradation characteristics while preserving aberration patterns, while FPD in Fig. 8(c) constructs aberration- and VG-dominant pseudo observations for separate LoRA branches. These results support $\mathbf{G}_{\text{tgt}}$ as a target cue for compound degradation adaptation.

5 CONCLUSION

In this paper, we address compound optical degradation involving co-existing aberration and veiling glare in simplified optical systems. We propose LENSBRIDGE, which adapts a reusable aberration correction foundation to target-specific compound degradation using only a few unpaired target observations, with a shared Fourier-amplitude statistic guiding both FDC and FPD. Experimental results across multiple optical systems demonstrate favorable restoration performance, effective few-shot adaptation, and applicability across different restoration backbones. We hope this work can facilitate efficient adaptation of reusable restoration foundations to diverse optical systems.

Limitations and Future Work. LENSBRIDGE relies on the coverage of the LensLib foundation, and broader target-specific aberration priors could further improve unseen lens adaptation. The current VG adaptation uses frequency-guided degradation completion rather than a fully calibrated physical model, motivating future exploration of broader optical degradation foundations and physically grounded VG modeling under complex illumination.

AI USE STATEMENT

In this work, we used generative AI tools to aid and polish the writing of the manuscript. We have reviewed and revised all AI-assisted text and verified the technical content and citations. We take responsibility for the final content of this work, including text and claims produced with the aid of generative AI.

REFERENCES

Chaofeng Chen, Xinyu Shi, Yipeng Qin, Xiaoming Li, Xiaoguang Han, Tao Yang, and Shihui Guo. Real-world blind super-resolution via feature matching with implicit high-resolution priors. In *MM*, 2022. 5

Qikai Chen, Jiacheng Zhou, Sijie Pian, Jingang Xu, Xingyi Li, Bihua Li, Chentao Lu, Zhuning Wang, Qi Jiang, Shanhe Qin, Hantao Zhan, Benhao Zhang, Xu Liu, Kaiwei Wang, and Yaoguang Ma. Hybrid meta-optics enabled compact augmented reality display with computational image reinforcement. *ACS Photonics*, 2024. 1

Shiqi Chen, Huajun Feng, Keming Gao, Zhihai Xu, and Yueting Chen. Extreme-quality computational imaging via degradation framework. In *ICCV*, 2021a. 1, 3

Shiqi Chen, Huajun Feng, Dexin Pan, Zhihai Xu, Qi Li, and Yueting Chen. Optical aberrations correction in postprocessing using imaging simulation. *ACM Transactions on Graphics*, 2021b. 1, 3

Shiqi Chen, Jingwen Zhou, Menghao Li, Yueting Chen, and Tingting Jiang. Mobile image restoration via prior quantization. *Pattern Recognition Letters*, 2023. 1

Xiang Chen, Jinshan Pan, Jiangxin Dong, Jian Yang, and Jinhui Tang. Foundir-v2: Optimizing pre-training data mixtures for image restoration foundation model. In *CVPR*, 2026. 8

Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz Khan. AdaIR: Adaptive all-in-one image restoration via frequency mining and modulation. In *ICLR*, 2025. 8

Yuekun Dai, Chongyi Li, Shangchen Zhou, Ruicheng Feng, and Chen Change Loy. Flare7K: A phenomenological nighttime flare removal dataset. In *NeurIPS*, 2022. 3

Yuekun Dai, Yihang Luo, Shangchen Zhou, Chongyi Li, and Chen Change Loy. Nighttime smartphone reflective flare removal using optical center symmetry prior. In *CVPR*, 2023. 3

Yuekun Dai, Chongyi Li, Shangchen Zhou, Ruicheng Feng, Yihang Luo, and Chen Change Loy. Flare7K++: Mixing synthetic and real datasets for nighttime flare removal and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 3, 8, 16

Thomas Eboli, Jean-Michel Morel, and Gabriele Facciolo. Fast two-step blind optical aberration correction. In *ECCV*, 2022. 8

Ruicheng Feng, Chongyi Li, Huaijin Chen, Shuai Li, Chen Change Loy, and Jinwei Gu. Removing diffraction image artifacts in under-display camera via dynamic skip connection networks. In *CVPR*, 2021. 18, 21

D. A. Fish, A. M. Brinicombe, E. R. Pike, and J. G. Walker. Blind deconvolution by means of the Richardson–Lucy algorithm. *Journal of the Optical Society of America A*, 1995. 3

Jin Gong, Runzhao Yang, Weihang Zhang, Jinli Suo, and Qionghai Dai. A physics-informed low-rank deep neural network for blind and universal lens aberration correction. In *CVPR*, 2024. 1

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *ICLR*, 2022. 4

Matthias Hullin, Elmar Eisemann, Hans-Peter Seidel, and Sungkil Lee. Physically-based real-time lens flare rendering. In *SIGGRAPH*, 2011. 2

Xiaozhong Ji, Guangpin Tao, Yun Cao, Ying Tai, Tong Lu, Chengjie Wang, Jilin Li, and Feiyue Huang. Frequency consistent adaptation for real world super resolution. In *AAAI*, 2021. 3

Qi Jiang, Shaohua Gao, Yao Gao, Kailun Yang, Zhonghua Yi, Hao Shi, Lei Sun, and Kaiwei Wang. Minimalist and high-quality panoramic imaging with PSF-aware transformers. *IEEE Transactions on Image Processing*, 2024. 3

Qi Jiang, Zhonghua Yi, Shaohua Gao, Yao Gao, Xiaolong Qian, Hao Shi, Lei Sun, JinXing Niu, Kaiwei Wang, Kailun Yang, and Jian Bai. Representing domain-mixing optical degradation for real-world computational aberration correction via vector quantization. *Optics & Laser Technology*, 2025. 1, 8, 16

Qi Jiang, Yao Gao, Shaohua Gao, Zhonghua Yi, Xiaolong Qian, Hao Shi, Kailun Yang, Lei Sun, Kaiwei Wang, and Jian Bai. OmniLens: Towards universal lens aberration correction via lenslib-to-specific domain adaptation. *Optics & Laser Technology*, 2026. 1, 2, 3, 8, 16

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. MUSIQ: Multi-scale image quality transformer. In *ICCV*, 2021. 7

Xiu Li, Jinli Suo, Weihang Zhang, Xin Yuan, and Qionghai Dai. Universal and flexible optical aberration correction using deep-prior based deconvolution. In *ICCV*, 2021a. 4

Zhaoyi Li, Peng Lin, Yao-Wei Huang, Joon-Suh Park, Wei Ting Chen, Zhujun Shi, Cheng-Wei Qiu, Ji-Xin Cheng, and Federico Capasso. Meta-optics achieves RGB-achromatic focusing for virtual reality. *Science Advances*, 2021b. 1

Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image restoration using swin transformer. In *ICCVW*, 2021. 8, 9, 16

Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *CVPRW*, pp. 1132–1140, 2017. 17

Chengxu Liu, Lu Qi, Jinshan Pan, Xueming Qian, and Ming-Hsuan Yang. Frequency domain-based diffusion model for unpaired image dehazing. In *ICCV*, 2025. 3

Jingren Liu, Shuning Xu, Qirui Yang, Yun Wang, Xiangyu Chen, and Zhong Ji. Fape-ir: Frequency-aware planning and execution framework for all-in-one image restoration. In *CVPR*, 2026. 8

Jun Luo, Yunfeng Nie, Wenqi Ren, Xiaochun Cao, and Ming-Hsuan Yang. Correcting optical aberration via depth-aware point spread functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024a. 3, 4

Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models for universal image restoration. In *ICLR*, 2024b. 8

Debabrata Mandal, Zhihan Peng, Yujie Wang, and Praneeth Chakravarthula. Enabling high-quality in-the-wild imaging from severely aberrated metalens bursts. *arXiv preprint arXiv:2510.10083*, 2025. 3

Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a "completely blind"' image quality analyzer. *IEEE Signal Processing Letters*, 2013. 7

Xiaolong Qian, Qi Jiang, Yao Gao, Lei Sun, Kailun Yang, Xian Wang, Zhonghua Yi, Wenyong Li, Ming-Hsuan Yang, Luc Van Gool, and Kaiwei Wang. Towards blind lens aberration correction via large lenslib pre-training and discrete degradation priors. In *ICCP*, 2026a. 1, 3, 5

Xiaolong Qian, Qi Jiang, Lei Sun, Zongxi Yu, Kailun Yang, Peixuan Wu, Jiacheng Zhou, Yao Gao, Yaoguang Ma, Ming-Hsuan Yang, and Kaiwei Wang. Learning latent transmission and glare maps for lens veiling glare removal. In *CVPR*, 2026b. 1, 2, 3, 6, 8, 16

Ramesh Raskar, Amit Agrawal, Cyrus A Wilson, and Ashok Veeraraghavan. Glare aware photography: 4D ray sampling for reducing glare effects of camera lenses. In *SIGGRAPH*, 2008. 2

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 4

Axel Sauer, Dominik Lorenz, Andreas Blattmann, and Robin Rombach. Adversarial diffusion distillation. In *ECCV*, 2024. 4, 6

Christian J. Schuler, Michael Hirsch, Stefan Harmeling, and Bernhard Schölkopf. Blind correction of optical aberrations. In *ECCV*, 2012. 3

Alexey Valeryevich Shoshin and Evgeny Alexandrovich Shvets. Veiling glare removal: synthetic dataset generation, metrics and neural network architecture. *Computer Optics*, 2021. 3, 4

Lei Sun, Xiaolong Qian, Qi Jiang, Xian Wang, Yao Gao, Kailun Yang, Kaiwei Wang, Radu Timofte, Danda Paudel, Luc Van Gool, Minjian Zhang, Liwen Zhang, Zhipei Lei, Kailing Tang, Zhe Xu, Xiang Li, Dingyong Gou, Yanlin Wu, Cong Li, Xiaohui Cui, Jianying Chan, Wenhao Miao, Yaoxin Jiang, Guoyi Xu, Jiajia Liu, Yaokun Shi, Jiachen Tu, Hao Yang, Ruikun Zhang, Liyuan Pan, Hiroto Shirono, Kosuke Shigematsu, Asuka Shin, Zixuan Zhou, Jingya Yang, Chunmiao Li, Jun Wang, Yimeng Yan, Ji Zhang, Jihye Park, Yoonjin Im, Chun Hyung-Ju, Hyunhee Park, MinKyu Park, and Wei Zhou. NTIRE 2026 the first challenge on blind computational aberration correction: Methods and results. In *CVPRW*, 2026. 6

Eino-Ville Talvala, Andrew Adams, Mark Horowitz, and Marc Levoy. Veiling glare in high dynamic range imaging. *ACM Transactions on Graphics*, 2007. 3

Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, Lei Zhang, Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, Kyoung Mu Lee, et al. NTIRE 2017 challenge on single image super-resolution: Methods and results. In *CVPRW*, 2017. 17

Ethan Tseng, Shane Colburn, James Whitehead, Luocheng Huang, Seung-Hwan Baek, Arka Majumdar, and Felix Heide. Neural nano-optics for high-quality thin lens imaging. *Nature Communications*, 2021. 6

Florin Vasluianu, Zongwei Wu, and Radu Timofte. SFNet - A spatial-frequency domain neural network for image lens flare removal. In *ICIP*, 2024. 3

Jianyi Wang, Kelvin C. K. Chan, and Chen Change Loy. Exploring CLIP for assessing the look and feel of images. In *AAAI*, 2023. 7

Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin C. K. Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *IJCV*, 2024. 6

Ruiyi Wang, Yushuo Zheng, Zicheng Zhang, Chunyi Li, Shuaicheng Liu, Guangtao Zhai, and Xiaohong Liu. Learning hazing to dehazing: Towards realistic haze generation for real-world image dehazing. In *CVPR*, 2025. 8, 16

Norbert Wiener. *Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications*. MIT press Cambridge, MA, 1949. 3

Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. Q-Align: Teaching LMMs for visual scoring via discrete text-defined levels. In *ICML*, 2024a. 7

Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. In *NIPS*, 2024b. 16

Yicheng Wu, Qiurui He, Tianfan Xue, Rahul Garg, Jiawen Chen, Ashok Veeraraghavan, and Jonathan T. Barron. How to train neural networks for flare removal. In *ICCV*, 2021. 3

Xinge Yang, Chuong Nguyen, Wenbin Wang, Kaizhang Kang, Wolfgang Heidrich, and Xiaoxing Li. Efficient depth-and spatially-varying image simulation for defocus deblur. In *ICCVW*, 2025. 4

Yanchao Yang and Stefano Soatto. FDA: Fourier domain adaptation for semantic segmentation. In *CVPR*, 2020. 3, 8

Hu Yu, Jie Huang, Yajing Liu, Qi Zhu, Man Zhou, and Feng Zhao. Source-free domain adaptation for real-world image dehazing. In *MM*, 2022. 3

Aiping Zhang, Zongsheng Yue, Renjing Pei, Wenqi Ren, and Xiaochun Cao. Degradation-guided one-step image super-resolution with diffusion priors. *arXiv preprint arXiv:2409.17058*, 2024. 5, 9, 16

Jianing Zhang, Jiayi Zhu, Feiyu Ji, Xiaokang Yang, and Xiaoyun Yuan. Degradation-modeled multipath diffusion for tunable metalens photography. *ICCV*, 2025. 3

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018a. 7

Zheng Zhang, Huajun Feng, Zhihai Xu, Qi Li, and Yueting Chen. Single image veiling glare removal. *Journal of Modern Optics*, 2018b. 3, 4

Jie Zhu and Sungkil Lee. PBFG: A new physically-based dataset and removal of lens flares and glares. In *ICCV*, 2025. 3

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 8

A Additional Implementation Details

Training Details. For the Stage-I foundation, training proceeds in two phases, with each phase optimized for 30K iterations. In Phase 1, we jointly optimize the PSF encoder, discrete codebook, attention-based condition mapping module, and aberration-aware LoRA branch using $\mathcal{L}_{\mathrm{rec}} + \mathcal{L}_{\mathrm{VQ}}$. This establishes the PSF-derived discrete aberration prior space. In Phase 2, we freeze the PSF encoder and codebook and optimize the image-based degradation predictor together with the restoration adaptation modules. The frozen PSF branch provides teacher supervision through $\mathcal{L}_{\mathrm{FM}}$, encouraging the image-based predictor to retrieve consistent entries from the established codebook. The Phase-2 objective is $\mathcal{L}_{\mathrm{rec}} + \mathcal{L}_{\mathrm{FM}}$. For target-system adaptation, we initialize from the Stage-I model and adapt each optical system independently using a limited number of unpaired target observations. Compared with the foundation training, Stage-II requires lightweight optimization with only 3K iterations for each target system. The pretrained SD-Turbo weights and discrete codebook remain frozen, while the image-based degradation predictor, attention-based condition mapping modules, VG encoder, and LoRA parameters of both branches are jointly optimized. The aberration branch is initialized from Stage I, while the VG branch is newly introduced. FDC-generated compound pairs are jointly sampled with the original LensLib aberration-only pairs during adaptation. Since the aberration-dominant observations generated by FPD are pseudo observations rather than physically exact aberration-only inputs, the feature matching loss is retained to preserve consistency with the Stage-I aberration prior space. The Stage-II objective combines the reconstruction loss and feature matching loss, i.e., $\mathcal{L}_{\mathrm{II}} = \mathcal{L}_{\mathrm{rec}} + \mathcal{L}_{\mathrm{FM}}$.

Target Adaptation Data. Tab. 5 summarizes the number of unpaired target observations used for Stage-II adaptation across different optical systems. For the SL and MRL systems, both `Screen` and `Real` domains are evaluated, while CAYE, Nano, and UDC provide additional real-world cases beyond the main benchmarks.

Table 5: Number of unpaired target observations used for few-shot Stage-II adaptation across different optical systems.

Optical System	Domain	# Target Observations	Purpose
SL	`Screen`	50	Main benchmark
SL	`Real`	50	Main benchmark
MRL	`Screen`	50	Main benchmark
MRL	`Real`	50	Main benchmark
CAYE	`Real`	26	Additional optical system
Nano	`Real`	3	Additional optical system
UDC	`Real`	30	OOD degradation evaluation

B Additional Analysis

Blind Prior Retrieval. We further analyze whether the image-based degradation predictor can retrieve PSF-derived priors from aberrated observations. Fig. 9 visualizes four representative lenses from the held-out LensLib split, including two single-element lenses, one smartphone lens, and one meta-optical system. For each lens, the cross denotes its PSF-derived prior, while the circles denote priors retrieved from different aberrated images under the same lens degradation. The retrieved priors concentrate around their corresponding PSF-derived anchors and remain separated across different optical systems, indicating that the predictor preserves lens-aware degradation information without requiring PSFs at inference.

Frequency-Domain Interpretation of Veiling Glare. For clarity, we omit the patch index p in the following derivation. From Eq. 1, the additional degradation introduced by veiling glare relative to the aberration-only observation can be written as

$$I_{\mathrm{de}} - I_{\mathrm{ab}} = I_{\mathrm{ab}} \odot (T - 1) + I_g. \tag{8}$$

Taking the Fourier transform gives

$$\mathcal{F}(I_{\mathrm{de}} - I_{\mathrm{ab}}) = \mathcal{F}(I_{\mathrm{ab}}) * \mathcal{F}(T - 1) + \mathcal{F}(I_g), \tag{9}$$

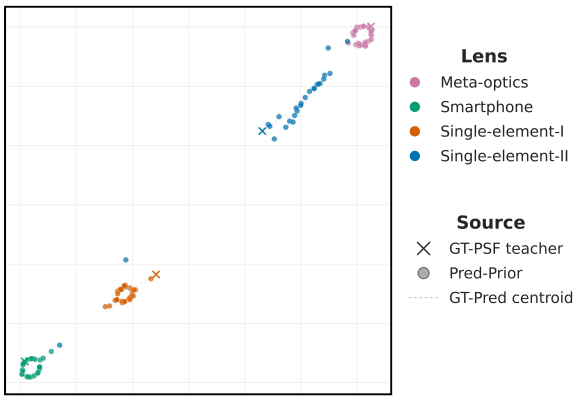

Figure 9: **t-SNE visualization of blind prior retrieval on representative held-out LensLib lenses.** Crosses denote PSF-derived priors, circles denote image-retrieved priors from aberrated observations, and gray lines connect matched GT–prediction pairs.

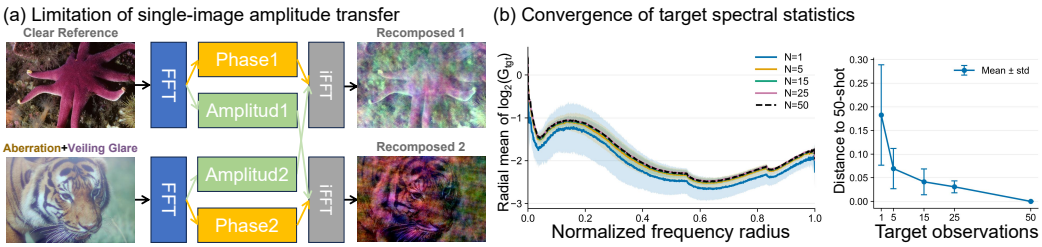

Figure 10: **Additional analysis of the target spectral statistic.** (a) Single-image amplitude transfer introduces content-dependent structures into the recomposed images. (b) The estimated target spectral statistic becomes progressively more stable as the number of target observations increases; the right plot reports the distance to the 50-shot statistic.

where $*$ denotes convolution in the frequency domain. The smooth diffuse component I_g naturally emphasizes low frequencies, whereas the multiplicative term couples the transmission variation with image content and is not strictly confined to the low-frequency region. Empirically, the introduction of veiling glare produces stronger low-frequency amplitude variations in our optical systems.

In contrast, lens aberration follows

$$I_{\text{ab}} = I_c \otimes K, \qquad \mathcal{F}(I_{\text{ab}}) = \mathcal{F}(I_c)\mathcal{F}(K), \tag{10}$$

where the optical transfer function $\mathcal{F}(K)$ affects image structures across different frequency ranges. Together with the observed spectral differences between aberration-only and compound degradations, this motivates the use of target-domain amplitude statistics in FDC and FPD.

Analysis of Target Spectral Statistics. We further analyze why the target statistic is estimated from multiple observations rather than directly transferring the Fourier amplitude of a single target image. As shown in Fig. 10(a), single-image amplitude transfer introduces scene-dependent structures into the recomposed images, since the Fourier amplitude contains both degradation and image-content information. Aggregating statistics over multiple target observations reduces such image-specific variations. In Fig. 10(b), the radial profiles progressively converge as more target observations are included, while the distance to the 50-shot statistic consistently decreases. These results support the use of an aggregated target-system spectral statistic for stable degradation completion.

Effect of Spatial Modulation. We further evaluate the spatial modulation mask used in FDC. As shown in Tab. 6, introducing spatial variation maintains comparable reconstruction fidelity while improving perceptual quality, particularly MUSIQ and Q-Align. This supports the use of spatial modulation to increase the diversity of the synthesized compound degradation.

Table 6: Effect of spatial modulation during FDC on `Screen-SL`.

Spatial Modulation	PSNR ↑	LPIPS ↓	MUSIQ ↑	Q-Align ↑
×	**20.505**	0.268	63.556	4.447
✓	20.462	**0.266**	**65.303**	**4.473**

Table 7: **Quantitative comparisons on the SL and MRL optical systems with additional target-system aberration-only paired supervision.** This setting supplements the primary setting with target-system aberration pairs. **Red** and blue indicate the best and the second-best performance, respectively.

Dataset	Metric	SwinIR	SwinIR+ Flare7K++	SwinIR+ DiffDehaze	DiffDehaze	QDMR	DeVeiler	OmniLens	LensBridge (Ours)
Screen SL	PSNR ↑	18.18	21.67	19.31	18.87	18.45	22.38	20.53	**22.69**
	SSIM ↑	0.686	0.723	0.642	0.657	0.681	**0.729**	0.642	0.679
	LPIPS ↓	0.298	0.297	0.347	0.435	0.291	0.261	0.262	**0.202**
	CLIP-IQA ↑	0.449	0.457	0.575	0.344	0.450	0.579	0.617	**0.636**
	NIQE ↓	5.400	3.996	5.873	7.843	4.262	4.163	**3.393**	3.428
	MUSIQ ↑	62.08	59.60	64.29	39.35	61.52	62.29	65.21	**68.20**
	Q-Align ↑	4.091	4.360	4.318	3.843	4.087	4.299	4.221	**4.709**
Screen MRL	PSNR ↑	19.34	20.74	18.28	18.55	20.67	**21.57**	20.64	21.46
	SSIM ↑	0.722	0.745	0.603	0.645	0.725	**0.746**	0.738	0.674
	LPIPS ↓	0.354	0.336	0.392	0.426	0.315	0.301	0.300	**0.266**
	CLIP-IQA ↑	0.458	0.469	0.550	0.460	0.527	0.526	0.569	**0.669**
	NIQE ↓	5.495	5.225	3.831	5.157	3.761	4.510	3.728	**3.263**
	MUSIQ ↑	58.72	58.22	64.70	54.96	61.29	58.43	62.84	**67.74**
	Q-Align ↑	4.206	4.356	4.375	4.343	4.344	4.310	4.435	**4.617**
Real SL	CLIP-IQA ↑	0.424	0.459	0.573	0.406	0.405	**0.607**	0.493	0.595
	NIQE ↓	5.710	5.616	6.122	6.467	4.770	4.446	4.692	**4.013**
	MUSIQ ↑	58.18	58.21	60.59	47.50	47.64	59.78	57.27	**63.89**
	Q-Align ↑	3.518	3.528	3.679	3.982	3.864	3.987	4.090	**4.451**
Real MRL	CLIP-IQA ↑	0.374	0.387	0.437	0.428	0.376	0.440	0.421	**0.633**
	NIQE ↓	6.696	6.345	5.226	6.840	5.478	5.286	5.505	**4.341**
	MUSIQ ↑	46.18	46.11	49.19	45.96	44.58	48.57	47.39	**62.87**
	Q-Align ↑	3.191	3.319	3.542	3.476	3.337	3.586	3.338	**4.045**

C More Results

C.1 Additional Quantitative Results

Complete Results under the Target Aberration-Paired Setting. We provide the complete quantitative results under the target aberration-paired setting in Tab. 7. The compared methods include cascaded pipelines combining SwinIR (Liang et al., 2021) with Flare7K++ (Dai et al., 2024) or DiffDehaze (Wang et al., 2025), as well as DiffDehaze, QDMR (Jiang et al., 2025), DeVeiler (Qian et al., 2026b), and OmniLens (Jiang et al., 2026). In addition to the paired `Screen` domain reported in the main paper, we include all metrics on both the `Screen` and `Real` domains.

Comparison with One-Step Diffusion Restoration Models. We further compare LensBridge with OSEDiff (Wu et al., 2024b) and S3Diff (Zhang et al., 2024), which also employ Stable-Diffusion-based one-step restoration backbones. Tab. 8 reports results under both the unpaired-target and target aberration-paired settings across the four SL/MRL domains. LensBridge achieves consistently favorable performance, particularly in perceptual and no-reference image quality metrics, providing a comparison with closely related one-step diffusion restoration models.

Target Data Efficiency. Tab. 9 provides the complete quantitative results corresponding to the few-shot analysis in the main paper. Using only five unpaired target observations improves PSNR from

Table 8: **Comparison with SD-based one-step diffusion restoration models under two target-adaptation settings.** Results are reported under the unpaired-target setting and the target aberration-paired setting. The best and second-best results within each setting are highlighted in **Red** and blue indicate the best and the second-best performance, respectively.

Dataset	Metric	Unpaired-Target			Target Aberration-Paired		
		OSEDiff	S3Diff	LensBridge (Ours)	OSEDiff	S3Diff	LensBridge (Ours)
Screen SL	PSNR ↑	**20.94**	20.07	20.46	22.39	21.72	**22.69**
	SSIM ↑	**0.677**	0.658	0.652	**0.679**	0.625	**0.679**
	LPIPS ↓	0.320	0.282	**0.266**	0.260	0.248	**0.202**
	CLIP-IQA ↑	0.524	0.593	**0.620**	0.581	**0.659**	0.636
	NIQE ↓	4.285	3.457	**3.417**	3.984	3.444	**3.428**
	MUSIQ ↑	60.68	**65.85**	65.30	66.05	68.13	**68.20**
	Q-Align ↑	4.289	4.440	**4.473**	4.554	**4.749**	4.709
Screen MRL	PSNR ↑	18.95	19.30	**19.72**	20.42	20.99	**21.46**
	SSIM ↑	**0.665**	0.659	0.652	**0.681**	0.664	0.674
	LPIPS ↓	0.351	0.335	**0.333**	0.306	0.298	**0.266**
	CLIP-IQA ↑	0.540	0.544	**0.556**	0.612	0.640	**0.669**
	NIQE ↓	3.648	3.514	**3.349**	3.366	3.482	**3.263**
	MUSIQ ↑	60.82	**61.08**	60.55	65.42	66.49	**67.74**
	Q-Align ↑	4.349	**4.488**	4.400	**4.657**	4.603	4.617
Real SL	CLIP-IQA ↑	0.507	0.523	**0.572**	0.566	**0.604**	0.595
	NIQE ↓	4.616	4.516	**4.150**	4.671	4.462	**4.013**
	MUSIQ ↑	57.62	58.26	**60.32**	62.35	63.81	**63.89**
	Q-Align ↑	4.247	4.334	**4.354**	4.289	4.416	**4.451**
Real MRL	CLIP-IQA ↑	0.533	0.515	**0.536**	0.567	0.571	**0.633**
	NIQE ↓	4.484	**4.201**	4.363	4.479	**4.253**	4.341
	MUSIQ ↑	55.31	55.57	**56.30**	60.24	60.42	**62.87**
	Q-Align ↑	3.875	3.907	**3.919**	**4.065**	4.009	4.045

Table 9: Few-shot adaptation with unpaired target observations on the Screen SL domain.

Target Images	PSNR↑	LPIPS↓	MUSIQ↑	Q-Align↑
0	15.075	0.332	53.651	3.725
5	19.551	0.277	64.283	4.384
15	19.654	0.271	64.579	4.420
25	20.034	0.270	65.113	4.436
50	**20.462**	**0.266**	**65.303**	**4.473**

15.075 to 19.551 and LPIPS from 0.332 to 0.277. Further increasing the number of target images yields smaller incremental improvements, consistent with the trend observed in the main paper.

Sensitivity to the Clear Reference Dataset. The target spectral statistic in Eq. 4 is computed using a generic clear-image collection. We evaluate the sensitivity to this choice on the Real-SL domain by replacing the default DIV2K collection (Timofte et al., 2017) with Flickr2K (Lim et al., 2017). As shown in Tab. 10, the results remain comparable between the two clear reference collections, indicating limited sensitivity to the choice of $\mathcal{D}_{\text{clr}}$.

C.2 Additional Visual Comparisons

Visual Comparisons under the Unpaired-Target Setting. Figs. 11 and 12 provide additional visual comparisons on the Screen and Real domains, respectively. Across both SL and MRL systems, LensBridge better recovers fine structures degraded by aberration while reducing the contrast and color degradation associated with veiling glare. The improvements remain consistent across controlled screen captures and in-the-wild observations.

Results on Additional Optical Systems. We provide additional results on the CAYE lens and nano-optics system, where the Stage-I foundation is independently adapted to each system using only a few unpaired degraded observations. As shown in Fig. 13, Stage I already alleviates aberration-

Table 10: Sensitivity to the clear reference dataset on the `Real` SL domain.

$\mathcal{D}_{\text{clr}}$	CLIP-IQA↑	NIQE↓	MANIQA↑	MUSIQ↑	Q-Align↑
Flickr2K	0.549	4.003	0.372	60.884	4.380
DIV2K	0.572	4.150	0.367	60.320	4.354

induced blur and recovers part of the image details, while subsequent target adaptation further improves contrast and color under compound degradation. The consistent behavior across both systems supports the applicability of the two-stage framework beyond the SL and MRL benchmarks.

Visual Comparisons under the Target Aberration-Paired Setting. We provide additional visual comparisons under the target aberration-paired setting in Figs. 14 and 15. Across both SL and MRL systems, LensBridge produces sharper structures and improved contrast compared with existing restoration and adaptation methods. These results complement the quantitative comparisons under the same setting, with consistent improvements across different optical systems.

Additional Results on Real UDC Images. We further evaluate LensBridge on real UDC images from Feng et al. (2021). UDC images exhibit severe contrast attenuation and haze-like appearance caused by complex optical artifacts, providing an additional out-of-distribution case beyond the current LensLib degradation distribution. As shown in Fig. 16, the method improves global contrast and visibility across diverse indoor and outdoor scenes, while fine diffraction artifacts remain challenging due to the absence of diffraction-aware PSFs in the current LensLib distribution. This experiment provides an additional qualitative evaluation of the proposed foundation-to-compound adaptation paradigm under unseen optical degradations.

Failure Cases and Future Work. Fig. 17 shows two representative challenging cases. For the SL system, fine high-frequency structures such as small text remain difficult to reconstruct accurately, partly due to latent-space compression and the absence of dedicated high-frequency restoration designs. For the MRL system, severe overexposure produces saturated regions with substantial information loss, preventing faithful recovery of the missing content. Future work will explore dedicated high-frequency restoration modules, more robust handling of severely saturated observations.

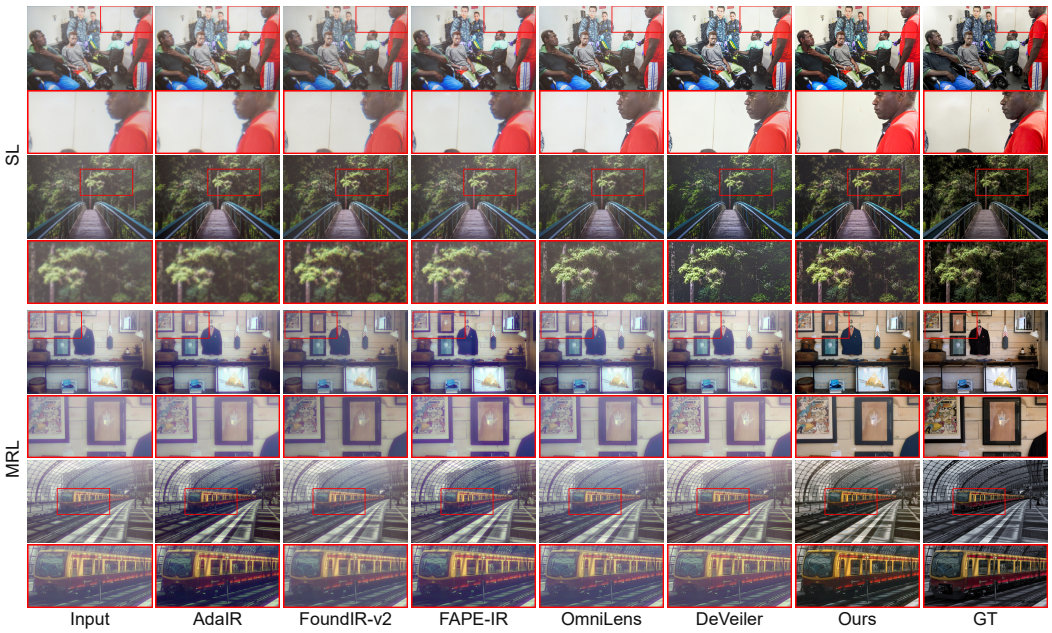

Figure 11: **Additional visual comparisons on the `Screen` domain under the unpaired-target setting.** Results are shown for both the SL and MRL optical systems. Zoom in for the best view.



Figure 12: **Additional visual comparisons on the `Real` domain under the unpaired-target setting.** Results are shown for both the SL and MRL optical systems. Zoom in for the best view.

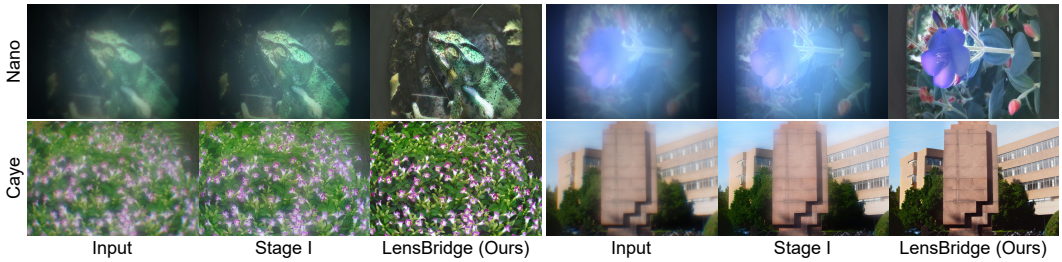

Figure 13: **Additional visual results on the CAYE lens and nano-optics system.** Stage I denotes the reusable aberration correction foundation before target adaptation, while LENSBRIDGE denotes the adapted model using unpaired target observations. Zoom in for the best view.

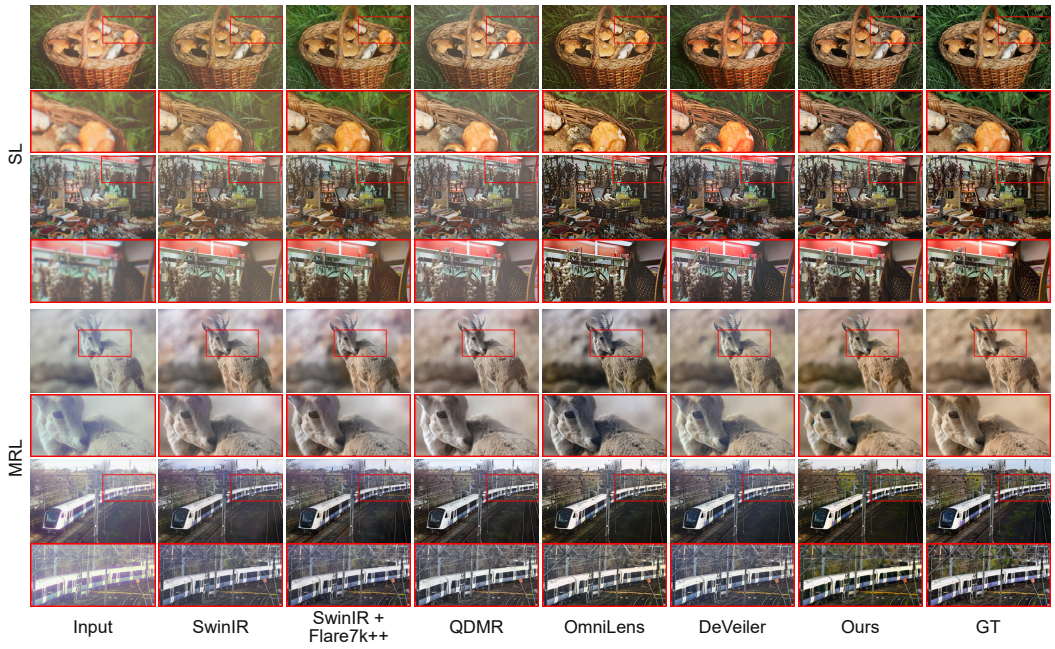

Figure 14: **Additional visual comparisons on the `Screen` domain under the target aberration-paired setting.** Results are shown for both the SL and MRL systems. Zoom in for the best view.



Figure 15: **Additional visual comparisons on the `Real` domain under the target aberration-paired setting.** Results are shown for both the SL and MRL systems. Zoom in for the best view.

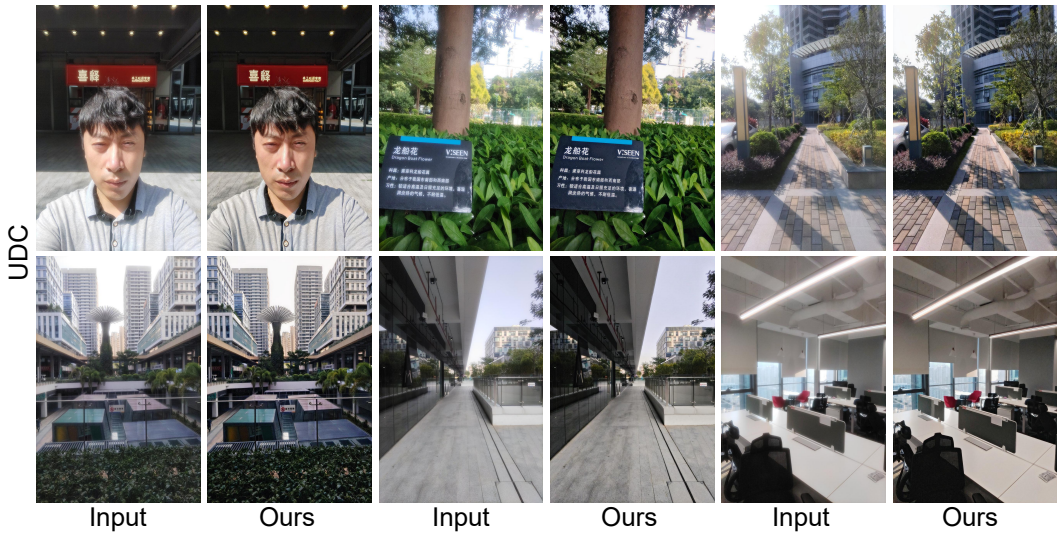

Figure 16: **Visual results on real UDC images** from Feng et al. (2021). LENSBRIDGE provides moderate restoration of contrast and visibility under UDC degradation. Zoom in for the best view.

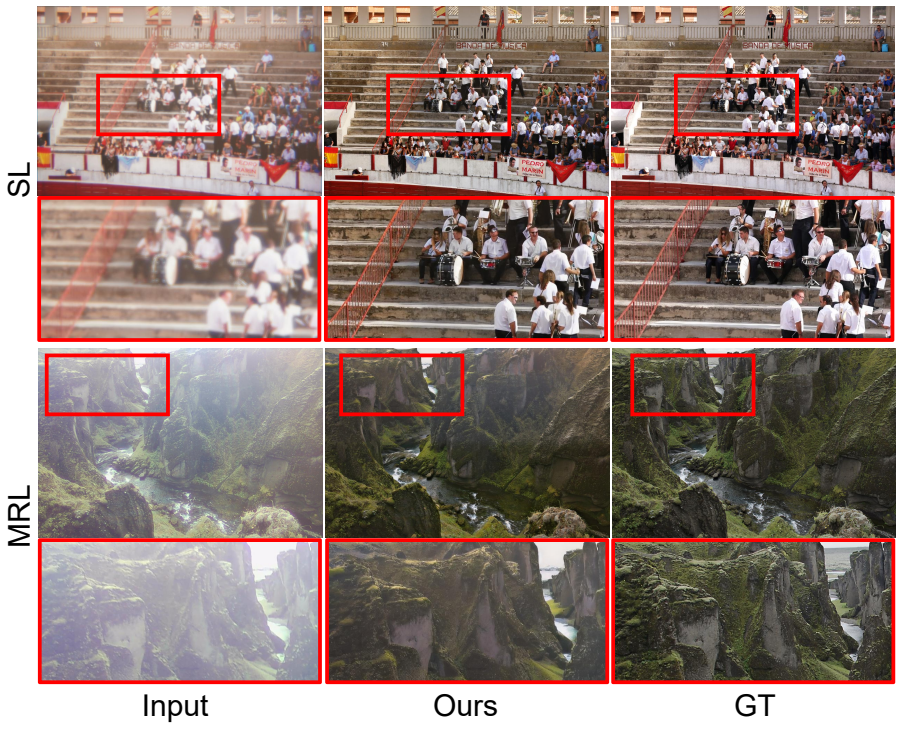

Figure 17: **Failure cases.** Fine high-frequency text in SL remains challenging to reconstruct faithfully, while severe overexposure in MRL causes substantial information loss.