

# DCA-LUT: Deep Chromatic Alignment with 5D LUT for Purple Fringing Removal

Jialang Lu<sup>1\*</sup>, Shuning Sun<sup>2\*</sup>, Pu Wang<sup>3</sup>, Chen Wu<sup>4</sup>, Feng Gao<sup>5</sup>, Lina Gong<sup>6</sup>, Dianjie Lu<sup>7</sup>, Guijuan Zhang<sup>7</sup>, Zhuoran Zheng<sup>8†</sup>

<sup>1</sup>Hubei University, Wuhan, China

<sup>2</sup>University of the Chinese Academy of Sciences, Beijing, China

<sup>3</sup>Shandong University, Jinan, China

<sup>4</sup>National University of Defense Technology, Changsha, China

<sup>5</sup>Ocean University of China, Qingdao, China

<sup>6</sup>Nanjing University of Aeronautics and Astronautics, Nanjing, China

<sup>7</sup>Shandong Normal University, Jinan, China

<sup>8</sup>Qilu University of Technology, Jinan, China

lujialang@stu.hubu.edu.cn, sunshuning23@mailsucas.ac.cn, wangou@mail.sdu.edu.cn, wuchen5X@mail.ustc.edu.cn, gaofeng@ouc.edu.cn, gonglina@nuaa.edu.cn, {ludianjie, zhangguijuan}@sdu.edu.cn, zhengzr@njjust.edu.cn

## Abstract

Purple fringing, a persistent artifact caused by Longitudinal Chromatic Aberration (LCA) in camera lenses, has long degraded the clarity and realism of digital imaging. Traditional solutions rely on complex and expensive apochromatic (APO) lens hardware and the extraction of handcrafted features, ignoring the data-driven approach. To fill this gap, we introduce DCA-LUT, the first deep learning framework for purple fringing removal. Inspired by the physical root of the problem, the spatial misalignment of RGB color channels due to lens dispersion, we introduce a novel Chromatic-Aware Coordinate Transformation (CA-CT) module, learning an image-adaptive color space to decouple and isolate fringing into a dedicated dimension. This targeted separation allows the network to learn a precise “purple fringe channel”, which then guides the accurate restoration of the luminance channel. The final color correction is performed by a learned 5D Look-Up Table (5D LUT), enabling efficient and powerful color mapping. To enable robust training and fair evaluation, we constructed a large-scale synthetic purple fringing dataset (PF-Synth). Extensive experiments in synthetic and real-world datasets demonstrate that our method achieves state-of-the-art performance in purple fringing removal.

**Extended version** — <https://arxiv.org/pdf/2511.12066>

## Introduction

In the pursuit of photorealism, digital imaging has become an indispensable technology. However, the quality of images is often compromised by optical artifacts inherent in lens systems. Among these, purple fringing is a persistent and visually disruptive artifact, manifesting as spurious purple or magenta halos around high-contrast edges. The physical

root of this phenomenon is Longitudinal Chromatic Aberration (LCA) (Clarke 2018), a simple camera lens refracts light of different wavelengths at slightly different angles, causing the focal planes of the primary colors (red, green, and blue) to separate (Volkova et al. 2019; Milster, Kim, and Wang 2022), illustrating in the Fig. 1. This spatial misalignment results in the characteristic fringing that degrades image sharpness and color fidelity. While hardware solutions like expensive apochromatic (APO) lenses can mitigate this issue (Ibrahim 2014; Sanli et al. 2023), they are often impractical for ubiquitous devices like smartphones. Consequently, developing effective computational methods for purple fringing removal remains a critical area of research.

Early computational approaches to this problem relied on handcrafted features to identify and correct purple fringing (Kim and Park 2008). This traditional approach begins by identifying high-contrast edges where fringing is most common. Subsequently, it desaturates the specific purple chromaticity, which is characterized by red and blue channel values being notably higher than the green one (Tomaselli et al. 2011). However, this rigid, rule-based logic is inherently unreliable. As shown in Fig. 2, it often mistakes legitimate purple objects for artifacts or fails to completely remove the fringe, leaving behind residual color casts (Kim and Park 2008; Malik and Karthik 2020). Furthermore, their performance is often inconsistent across different lens types and lighting conditions, limiting their general applicability.

Neural networks powered by deep learning have led to significant advances in computational image restoration tasks such as denoising, deblurring, and super-resolution (Jiang et al. 2025; Zamir et al. 2022; Liang et al. 2021). This success has inspired initial efforts to apply deep learning to correct general optical aberrations, including chromatic aberration (Maman et al. 2023). Concurrently, learnable Look-Up Tables (LUTs) have emerged as a highly efficient method for performing complex, non-linear color

\*These authors contributed equally.

†Corresponding author.

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

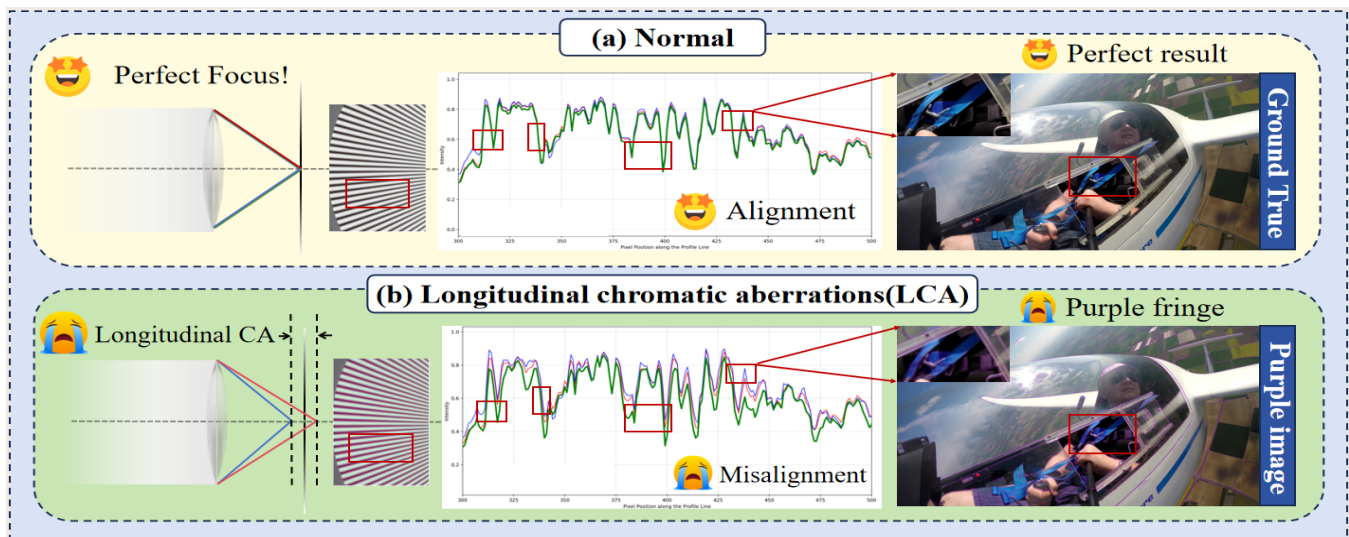


Figure 1: This figure illustrates the physical root of Longitudinal Chromatic Aberration (LCA). (a) Normal Case: All color channels are perfectly aligned, resulting in a sharp, clean image. (b) Case with LCA: Different colors focus at different depths, causing red/blue shift relative to green, resulting in purple fringing.

transformations in image enhancement (Zeng et al. 2020). However, their application has been predominantly limited to global tasks like color grading, rather than precise, local artifact correction. We identify a critical opportunity to adapt this powerful and efficient technology to tackle the purple fringing problem directly.

Motivated by the physics of LCA, we introduce DCA-LUT to fill this gap, a framework that first decouples the fringe artifact into a dedicated channel, then uses that channel’s geometric gradient to guide a direction-aware 5D LUT for precise, targeted correction. The main contributions of our work can be summarized as follows:

- We propose a Chromatic-Aware Coordinate Transformation (CA-CT) module inspired by the physics of LCA. This module learns to map the misaligned RGB channels into a new space, cleanly separating the purple fringe aberration into a dedicated channel.
- We design a high-dimensional, gradient-guided 5D LUT module for chromatic aberration correction. It performs precise and efficient non-linear color correction. This design is targeted for purple fringing removal, where directional fringe cues are crucial for accurate correction.
- To facilitate robust training and standardized evaluation, we construct **PF-Synth**, a large-scale synthetic dataset for purple fringing. We also propose the **Edge Chromatic Aberration Score (ECAS)**, a task-specific metric that quantifies residual color fringing along high-contrast edges, providing more accurate evaluation than PSNR and SSIM. Extensive experiments demonstrate that DCA-LUT achieves state-of-the-art performance on both synthetic and real-world images.




Figure 2: Comparison of fringe detection between traditional heuristic methods (Kim and Park 2010) and our learned CA-CT module. Our method accurately isolates the fringe artifact while preserving legitimate purple content.

## Related work.

**Chromatic Aberration Correction.** Computational approaches to purple fringing correction have evolved from heuristic-based models to general deep learning frameworks. Early methods focused on detecting fringe regions using handcrafted features, such as edge information and color gradients, and then applied local corrections like color channel shifting or targeted desaturation (Kang 2007; Kim and Park 2010). However, as shown in Fig. 2, these traditional techniques often struggle to generalize across diverse scenes and can introduce visual artifacts. With the rise of deep learning, Convolutional Neural Networks have been successfully applied to correct a broad range of optical aberrations by treating the problem as a blind deconvolution or geometric warping task (Schuler et al. 2015; Li et al. 2019, 2021). While powerful, these general-purpose networks typically lack data-driven and physical inductive bias for the unique channel-specific defocus that characterizes Longitu-

dinal Chromatic Aberration (LCA), failing to specifically model the root cause of purple fringing. Our work addresses this gap by introducing a framework specifically engineered to isolate and correct purple fringing.

**Learnable Look-Up Tables for Image Enhancement.** Integrating Look-Up Tables (LUTs) into deep learning pipelines has emerged as a highly efficient method for performing complex, non-linear color transformations. This was pioneered by learning a 3D LUT for real-time photo enhancement, a concept that was later extended to improve flexibility and reduce memory through innovations like separable structures and adaptive color space sampling (Zeng et al. 2020; Yang et al. 2022b,a). To date, the application of learnable LUTs has been predominantly limited to global image enhancement tasks, such as color grading and style transfer (Conde et al. 2024). In contrast, we repurpose the LUT as a precise, local artifact correction module. We achieve this by introducing a novel direction-aware 5D LUT that utilizes cross-channel guidance, informing the geometric gradient of the isolated purple fringe channel into the LUT’s coordinates to steer the final correction. Motivated by the physics of LCA, our direction-aware LUT directly counteracts the local color channel misalignment, enabling targeted, orientation-sensitive corrections that efficiently eliminate residual color shifts on resource-constrained devices.

## Methodology

Our proposed framework, DCA-LUT, is engineered to remove purple fringing by first transforming the input image into a learned color space. Within this space, the isolated luminance and fringe channels are independently corrected by LUT. The entire process is trained end-to-end, guided by a composite loss function that ensures both reconstruction fidelity and the physical plausibility of the learned transformations. The overall architecture is illustrated in Fig. 3.

### Chromatic-Aware Transformation

To isolate purple fringing, the first stage of our methodology generates image-specific parameters for transforming the input RGB image  $\mathbf{I}_{RGB}$  into our proposed Chromatic Aberration Space (CAS). This is achieved via an image-adaptive  $3 \times 3$  matrix  $\mathbf{M}$ , predicted by our Chromatic-Aware Coordinate Transformation (CA-CT) module. This module utilizes a lightweight ConvNeXt-based encoder to regress the elements of  $\mathbf{M}$  from the global features of the input image. The transformation is a per-pixel matrix multiplication. For each pixel  $p$  with color vector  $\mathbf{I}_{RGB}(p)$ , the corresponding pixel in the Chromatic Aberration Space is computed as:

$$\mathbf{I}_{CAS}(p) = \mathbf{M} \cdot \mathbf{I}_{RGB}(p). \quad (1)$$

This transformed pixel  $\mathbf{I}_{CAS}(p)$  is composed of a luminance channel ( $C_{lum}$ ), a fringe channel ( $C_{fringe}$ ), and an orthogonal channel ( $C_{ortho}$ ). To ensure the stability and physical plausibility of this learned transformation, we introduce an **axis alignment loss**,  $L_{align}$ , which encourages the basis vectors of the CAS to be mutually orthogonal:

$$L_{align} = \frac{1}{B} \sum_{b=1}^B \sum_{i \neq j, i, j \in \{1, 2, 3\}} (\mathbf{r}_{b,i} \cdot \mathbf{r}_{b,j})^2, \quad (2)$$

where  $\mathbf{r}_{b,i}$  is the  $i$ -th row vector of  $\mathbf{M}$  for the  $b$ -th image in a batch of size  $B$ .

### LUT Correction in Chromatic Aberration Space

Once the image is represented in CAS, we introduce our core innovation: a direction-aware correction mechanism that operates on the decoupled channels. This process is motivated by the physics of LCA; since the artifact is a spatial misalignment of color, its correction must be spatially informed. To this end, we employ a novel cross-channel guidance strategy, where the geometric properties of the fringe channel are used to steer the correction of the luminance channel.

This strategy is embodied in our **Direction-Aware 5D Look-Up Table (LUT)**, designed to adjust the luminance channel ( $C_{lum}$ ) by incorporating spatial context to prevent such artifacts. Instead of being a simple color mapper, its lookup coordinates are designed to fuse spatial context with chromatic directional cues. For each pixel  $p$ , the five dimensions for LUT lookup include the luminance values of the center pixel  $C_{lum}(p)$ , the pixel to the left  $C_{lum}(p_{left})$ , and the pixel above  $C_{lum}(p_{top})$ , where non-symmetric, L-shaped sampling effectively captures edge information, along with the horizontal  $\nabla_x C_{fringe}(p)$  ( $dx$ ) and vertical  $\nabla_y C_{fringe}(p)$  ( $dy$ ) gradients of the fringe channel, enabling encode the direction and magnitude of the purple fringe at pixel. By conditioning the luminance correction on the fringe’s gradient, the 5D-LUT learns to locally counteract the color channel misalignment at the heart of LCA, especially around high-frequency areas where fringing occurs, yielding the corrected luminance channel  $C'_{lum}$ .

Concurrently, the fringe channel ( $C_{fringe}$ ) is refined by a simpler, learnable **1D LUT**. This LUT performs a global mapping to suppress the residual fringe intensity across the entire image, producing the corrected fringe channel  $C'_{fringe}$ . To ensure the learned transformations are smooth and avoid introducing new artifacts, we apply a **smoothness loss**,  $L_{smooth}$ , which penalizes large gradients between adjacent entries in the LUTs. It is defined as:

$$L_{smooth} = \|\nabla^2 \mathbf{L}_{1D}\|_F^2 + \sum_{d=1}^5 \|\nabla_d \mathbf{L}_{5D}\|_F^2, \quad (3)$$

where  $\nabla^2$  and  $\nabla_d$  are the second-order and first-order finite difference operators, respectively. After correction, the corrected channels  $C'_{lum}$  and  $C'_{fringe}$ , along with the original orthogonal channel  $C_{ortho}$ , form the final corrected CAS image. This image is then transformed back to the RGB domain via the analytical inverse of the transformation matrix,  $\mathbf{M}^{-1}$ , to yield the final output image  $\mathbf{I}_{out}$ :

$$\mathbf{I}_{out} = \mathbf{M}^{-1} \cdot \begin{pmatrix} C'_{lum} \\ C'_{fringe} \\ C_{ortho} \end{pmatrix}. \quad (4)$$

The fidelity of this output is governed by a comprehensive **reconstruction loss**,  $L_{rec}$ . This loss is a weighted sum of a standard L1 pixel loss ( $L_{L1}$ ) and our **YCbCr Perceptual Loss** ( $L_p$ ):

$$L_{rec} = \lambda_{L1} L_{L1} + \lambda_p L_p. \quad (5)$$

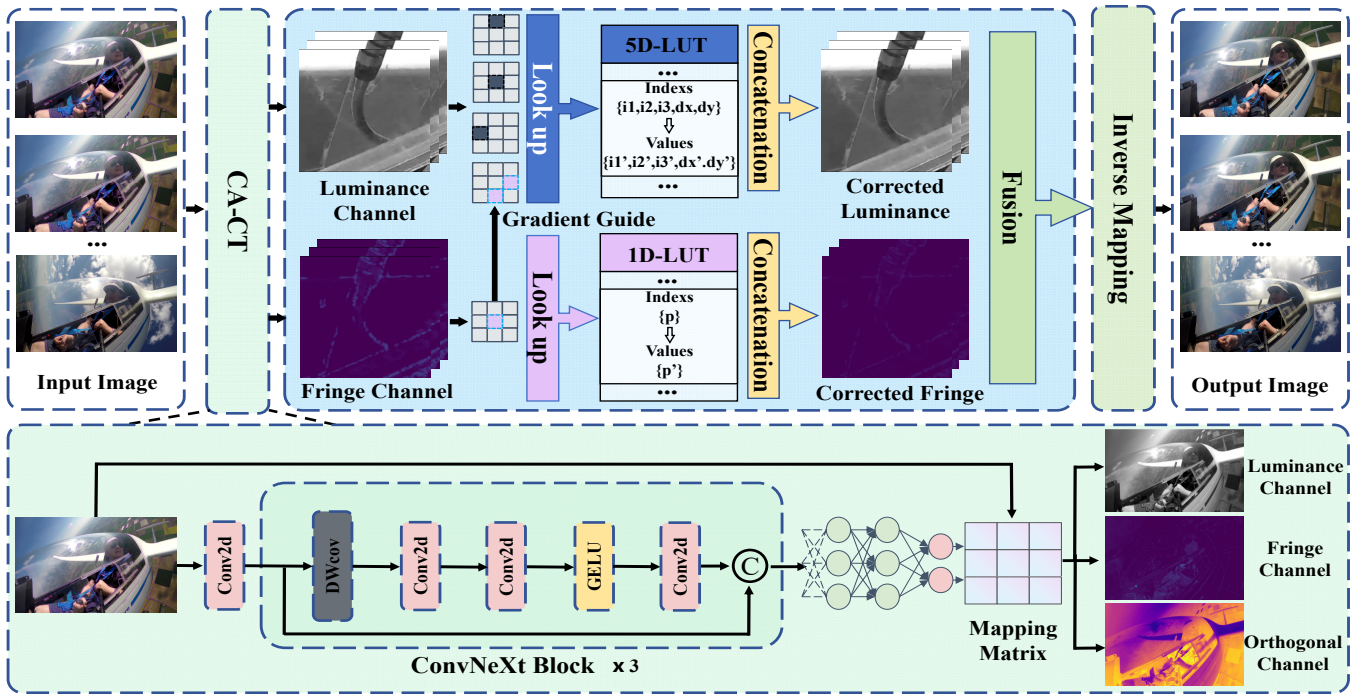


Figure 3: **The overall pipeline of DCA-LUT framework.** (a) **Chromatic-Aware Transformation (CA-CT).** The input image is first processed by a ConvNeXt-based encoder to map into our learned Chromatic Aberration Space, decoupling the image into Luminance Channel ( $C_{lum}$ ), Fringe Channel ( $C_{fringe}$ ), and an Orthogonal Channel. (b) **Direction-Aware Correction:** A dual-LUT correction branch, where the luminance channel is corrected by Direction-Aware 5D-LUT, utilizing the geometric gradients from the fringe channel to guide the luminance correction. Concurrently, the fringe channel is refined by a 1D-LUT.

The  $L_p$  term operates in the YCbCr color space. It computes a perceptual loss on the luma (Y) channel by measuring the L1 distance between deep features extracted by a pre-trained VGG19 network ( $\phi$ ) (Simonyan and Zisserman 2014). Simultaneously, it computes a direct L1 loss on the chroma (Cb, Cr) channels. This hybrid approach ensures both structural/textural fidelity on the luminance and color accuracy on the chrominance. It is defined as:

$$L_p = \|\phi(\mathbf{Y}_{out}) - \phi(\mathbf{Y}_{gt})\|_1 + \lambda_{chroma} (\|\mathbf{Cb}_{out} - \mathbf{Cb}_{gt}\|_1 + \|\mathbf{Cr}_{out} - \mathbf{Cr}_{gt}\|_1). \quad (6)$$

### Total Objective

The entire DCA-LUT framework is optimized end-to-end by minimizing a composite objective function. As implemented in the provided code, this function is a weighted sum of the reconstruction and regularization loss components:

$$L_{total} = L_{rec} + \lambda_1 L_{smooth} + \lambda_2 L_{align}. \quad (7)$$

This comprehensive objective function ensures that the network not only produces a high-fidelity reconstruction but also learns a physically plausible, stable, and regularized internal mechanism for purple fringing removal.

## Experiments

### Setup

Our framework is implemented using PyTorch and trained on a single NVIDIA RTX 3090 GPU. We use the Adam optimizer with a weight decay of  $5 \times 10^{-4}$ . The initial learning rate is set to  $5 \times 10^{-5}$  and decays using a Cosine Annealing schedule over 150 epochs. The batch size is set to 16.

### Datasets

The absence of large-scale, paired datasets for purple fringing removal necessitates the creation of a synthetic dataset for robust training and quantitative evaluation. To this end, we introduce **PF-Synth**, a new synthetic dataset constructed using high-quality video frames from the DAVIS 2017 and 2019 datasets (Pont-Tuset et al. 2017; Caelles et al. 2019), where frames from DAVIS 2017 are used for training, while those from DAVIS 2019 serve as the test and validation sets. The synthesis pipeline, detailed in the Appendix, is designed to mimic the physical properties of Longitudinal Chromatic Aberration. Given a clean image, we first detect its high-contrast edges using a Canny edge detector. A random subset of these edge pixels is then selected to ensure the fringing appears sparsely, as it does in reality. A mask is generated by dilating these sparse edge points, which is then blurred to create smooth transitions. Finally, this mask is used to blend a predefined purple color with the original image, with randomized parameters for intensity, width, and sparsity to

Method	Venue	Quality Metrics					Efficiency Metrics		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	ECAS $\downarrow$	$\Delta E$ $\downarrow$	Params (M) $\downarrow$	Time (s) 2K $\downarrow$	Time (s) 4K $\downarrow$
SRLUT	CVPR'21	38.376	0.9624	0.0411	0.0632	0.9581	0.15	0.0402	0.1768
SwinIR	ICCV'21	36.486	0.7923	0.1232	0.1755	2.3823	11.50	OOM	OOM
Restormer	CVPR'22	37.856	0.8896	0.0879	0.0662	2.1612	26.11	OOM	OOM
DDRM	NeurIPS'22	<b>39.154</b>	<u>0.9734</u>	<u>0.0317</u>	<u>0.0411</u>	<u>0.9581</u>	72.52	OOM	OOM
Lightendiffusion	ECCV'24	37.765	0.8724	0.0672	0.0715	0.9241	27.83	0.9624	OOM
NILUT	AAAI'24	38.445	0.9681	0.0416	0.0686	1.4218	<b>0.03</b>	0.0549	0.2794
DnLUT	CVPR'25	38.902	0.9642	0.0627	0.0831	1.5526	0.59	4.7505	OOM
BPAM	ICCV'25	38.165	0.9671	0.0588	0.1085	1.5755	11.68	0.4851	OOM
<b>Ours</b>	-	<u>39.052</u>	<b>0.9849</b>	<b>0.0303</b>	<b>0.0438</b>	<b>0.6784</b>	<u>0.13</u>	<b>0.0405</b>	<b>0.1756</b>

Table 1: Quantitative comparison with state-of-the-art methods on the purple fringing benchmark. For PSNR, SSIM, higher is better ( $\uparrow$ ). For LPIPS, ECAS,  $\Delta E$ , Parameters, and Inference Time, lower is better ( $\downarrow$ ). The best results are in **bold**, and the second-best are underlined. OOM stands for ‘‘Out of Memory’’, indicating that the method failed to process the corresponding resolution due to insufficient memory.

ensure diversity. For comprehensive evaluation, we also collected real-world purple fringe photos from photography forums to assess the model’s practical performance in real scenarios, with various consumer-grade lenses exhibiting noticeable purple fringing.

## Evaluation Metrics

**Standard Metrics.** To assess the overall image restoration quality, we employ three widely-used metrics: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM), which measure pixel-level fidelity and structural preservation, respectively. Additionally, we use the Learned Perceptual Image Patch Similarity (LPIPS), which better aligns with human perceptual judgment by comparing deep features. We also utilize the CIEDE2000 color difference metric  $\Delta E$  for quantifying the perceptual difference between two colors.

**Specialized Metric for Chromatic Aberration.** General metrics like PSNR and SSIM, while useful, are often insensitive to localized and color-specific artifacts such as purple fringing. A high PSNR score does not guarantee the complete removal of subtle color fringes along high-contrast edges. To address this limitation and provide a more accurate evaluation for this specific task, we propose a novel metric: the **Edge Chromatic Aberration Score (ECAS)**. The core idea of ECAS is to first identify high-contrast edge regions (where fringing occurs) using the ground-truth image, and then quantify the proportion of pixels within these regions that exhibit undesirable coloration in the corrected image. A lower ECAS score indicates a more effective suppression of chromatic aberration. Specifically, for a given corrected image  $\mathbf{I}_{corr}$  and ground-truth  $\mathbf{I}_{gt}$ , the calculation involves three steps. First, an edge mask  $\mathcal{M}_{edge}$  is generated from  $\mathbf{I}_{gt}$  by applying a Sobel filter and thresholding the gradient magnitude at  $\tau_{edge}$ . This mask identifies all high-contrast regions of interest. Second, an aberration mask  $\mathcal{M}_{aberration}$  is created from  $\mathbf{I}_{corr}$  by converting it to the HSV color space and identifying pixels whose hue falls within predefined purple/green ranges and whose saturation exceeds a threshold  $\tau_{sat}$ .

The final ECAS score is the ratio of the fringe pixels detected on edges to the total number of edge pixels:

$$ECAS = \frac{\sum(\mathcal{M}_{edge} \odot \mathcal{M}_{aberration})}{\sum \mathcal{M}_{edge}}, \quad (8)$$

where  $\odot$  denotes element-wise multiplication. A lower ECAS score indicates a more effective suppression of chromatic aberration, with a score of 0 being a perfect removal. This targeted metric serves as a crucial tool for evaluating model performance and validating our design choices.

## Quantitative and Qualitative Evaluation

We conduct a comprehensive comparison against a range of state-of-the-art (SOTA) methods, spanning different architectures and publication years. The competitors include LUT-based models (SRLUT (Jo and Kim 2021), NILUT (Conde et al. 2024), DNLUT (Yang et al. 2025), BPAM (Lou et al. 2025)), powerful Transformer-based general image restorers (SwinIR (Liang et al. 2021), Restormer (Zamir et al. 2022)), and recent diffusion-based models (Lightendiffusion (Jiang et al. 2024), DDRM (Kawar et al. 2022)).

**Quantitative Comparison.** Table 1 presents the quantitative results on our purple fringing benchmark. The results clearly demonstrate the superiority of our proposed method. Our model achieves state-of-the-art results in several key quality metrics, securing the best SSIM (0.9849), LPIPS (0.0303), and color accuracy with the lowest  $\Delta E$  (0.6784). While the diffusion-based model DDRM achieves the highest PSNR and lowest ECAS, our method is a very close second in both categories, confirming its top-tier reconstruction fidelity and artifact removal capability. Notably, our ECAS score of 0.0438 is significantly better than those of transformers like SwinIR and Restormer, underscoring the effectiveness of our specialized, physics-informed design.

**Qualitative Comparison.** As shown in the visual comparisons in Fig. 4, our method produces qualitatively superior results. General-purpose models like SwinIR, Restormer, despite their power, tend to either leave residual purple

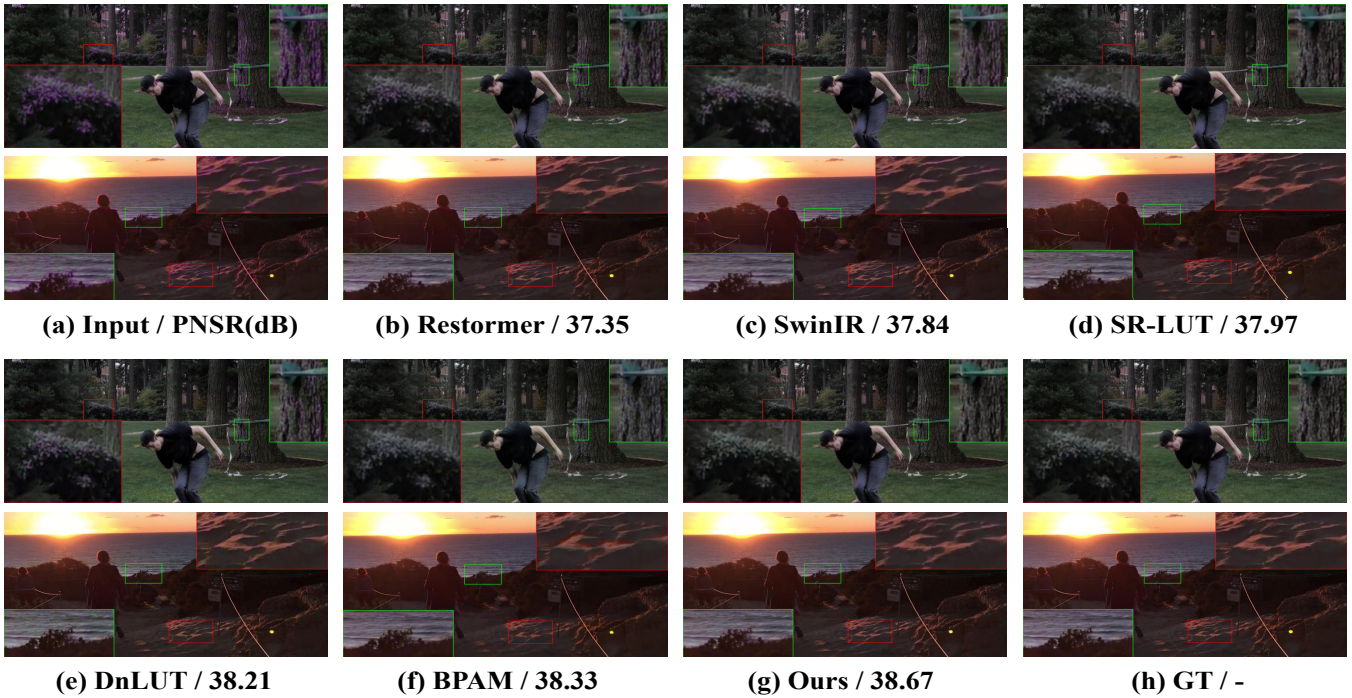


Figure 4: Visual comparison on our PF-Synth dataset. While competing methods often introduce secondary artifacts such as **darkening**, **color casts**, or **detail loss**, our approach is the only one that cleanly removes the purple fringe while faithfully preserving the original image content. (Please zoom in for the best view.)

Configuration	PSNR $\uparrow$	ECAS $\downarrow$
Correct Green Channel	38.104	0.0957
Correct Luma Channel (Ours)	<b>39.052</b>	<b>0.0438</b>
Symmetric LUT	37.552	0.1512
+ Asymmetric Sampling	38.827	0.0865
+ Gradient Guide (Ours)	<b>39.052</b>	<b>0.0438</b>

Table 2: Comprehensive ablation studies validating our key design choices.

fringe or over-smooth fine details along the edges. Specialized LUT-based methods also exhibit distinct drawbacks. Both SR-LUT and DnLUT manage to suppress the purple color but do so at the cost of creating unnatural dark outlines that degrade local contrast. BPAM fails more catastrophically by replacing the purple fringe with severe yellowish and greenish color artifacts, demonstrating a fundamental lack of robustness. In contrast, our method cleanly removes the chromatic aberration while faithfully preserving the underlying scene. Our results are visually the most plausible and closest to the ground truth, proving that our approach can precisely eliminate the artifact without causing collateral damage to the image content.

### Ablation Studies

To validate the critical components of our DCA-LUT framework, we conduct a series of ablation experiments. We sys-

tematically analyze two key design choices: the selection of the luminance channel for correction, and the architecture of our Direction-Aware 5D LUT.

**Advantages of Luminance-Based Correction over Green Reference.** Given that LCA is a misalignment of the red and blue channels relative to the green channel, a plausible alternative is to use the green channel as the reference for correction. However, this approach has a significant drawback. As illustrated in Fig. 5, directly correcting the image based on the green channel introduces a noticeable greenish color cast across the entire image, altering the global color balance. By using our CA-CT module to generate a dedicated luminance channel, our framework can perform corrections in a color-neutral space. This ensures that the artifact is removed without introducing undesirable color shifts, leading to a more natural and physically accurate result.

**Effectiveness of the Direction-Aware 5D LUT.** We validate the design of our core correction module, the Direction-Aware 5D LUT. We compare our design against two simpler variants to demonstrate the importance of both its asymmetric structure and the cross-channel gradient guidance. The qualitative results in Fig. 6 and quantitative metrics in Table 2 clearly show the progression. The Symmetric LUT struggles to precisely target the fringe and the Asymmetric (L-shaped) LUT performs better by capturing the edge-like nature of the artifact more effectively. However, our final Gradient-Guided LUT achieves the best results. By explicitly using the fringe’s direction and magnitude as a guide, it performs the most precise correction, completely removing

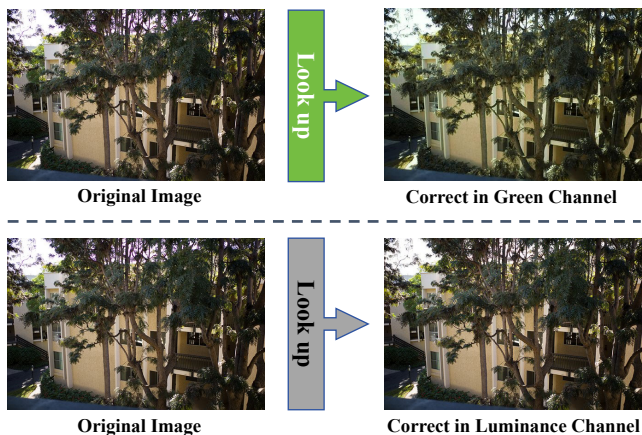


Figure 5: Comparison of correction strategies. Green-channel correction (top) introduces a noticeable green tint, while our luminance-based approach (bottom) preserves natural color fidelity.

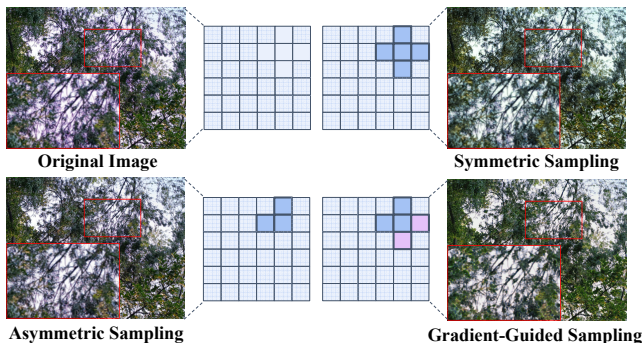


Figure 6: Visual ablation of the 5D LUT design. The gradient-guided approach (bottom right) most effectively removes the purple fringe without introducing artifacts, demonstrating the superiority of the direction-aware design.

the artifact while preserving the sharpness and integrity of the underlying edge.

## Discussion

**Architectural Rationale: 5D LUT for Structural Repair, 1D LUT for Global Color.** Purple fringing is edge-localized and structure-dependent. Our 5D LUT, in the Chromatic Aberration space with luminance/fringe cues and edge geometry ( $|\nabla|$ ), performs direction-aware local corrections on luminance, removing fringes while preserving detail. Global color fidelity is direction-invariant; a lightweight 1D LUT serves as the final tone curve, is easily regularized (with monotonicity constraints), and avoids direction-dependent casts at negligible cost. This separation lets 5D fix the structure and 1D enforce global color consistency. As shown in Table 1, our architecture delivers state-of-the-art performance with minimal computational cost.

**Real-world Purple Fringing Removal.** We validate our model on challenging real-world images, which often exhibit more complex fringing than synthetic data. As shown

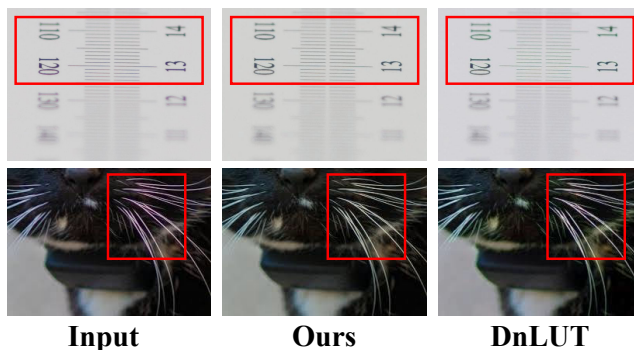


Figure 7: Visual comparison on real-world images. The competing method introduces a noticeable **greenish cast** and leaves **residual purple fringes**, while our method removes the fringe cleanly and preserves natural brightness and texture.

in Fig. 7, both our method and the competing specialized method DnLUT can suppress the purple hue. However, DnLUT introduces a significant darkening artifact, creating an unnatural dark halo around the high-contrast edges of the ruler’s markings and the cat’s whiskers. In contrast, our method cleanly eliminates the fringe while faithfully preserving the original brightness and texture of the scene. This demonstrates our model’s superior robustness and fidelity in practical scenarios.

**Model Complexity and Efficiency.** Our model excels in efficiency. For evaluation, we report the number of parameters (Params) and the average inference time on a single NVIDIA RTX 3090 GPU for 2K and 4K resolution images. As shown in Table 1, our method is the second most lightweight model, with only 0.13M parameters, and achieves the fastest inference speeds. This stands in stark contrast to large transformer and diffusion-based models (e.g., SwinIR, DDRM), which fail to process high-resolution inputs due to memory constraints. Overall, our method achieves an excellent trade-off between performance and efficiency.

## Conclusion

In this paper, we presented DCA-LUT, a novel and efficient framework for purple fringing removal. Motivated by the physics of Longitudinal Chromatic Aberration, our method first learns a transformation to isolate the fringe artifact, and then uses a Direction-Aware 5D LUT that leverages the fringe’s geometric gradient to steer the luminance restoration. Extensive experiments show that DCA-LUT achieves outstanding results, cleanly removing artifacts without the side effects seen in competing methods, all while being orders of magnitude faster and more memory-efficient. Moreover, we validated our model’s robustness through real-world tests across diverse scenarios and deployed it on mobile devices. Our work not only provides a practical solution to a persistent optical problem but also demonstrates a new paradigm for using learnable LUTs as precise, spatially-aware correction operators.

## Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grant 62172265, in part by the Shandong Provincial Natural Science Foundation under Grant ZR2025MS1025.

## References

- Caelles, S.; Pont-Tuset, J.; Perazzi, F.; Montes, A.; Maninis, K.-K.; and Van Gool, L. 2019. The 2019 DAVIS Challenge on VOS: Unsupervised Multi-Object Segmentation. *arXiv:1905.00737*.
- Clarke, T. 2018. Chromatic Aberration in Digital Photomicrographs from Microscopes Requiring Compensating Eyepieces. *Microscopy Today*.
- Conde, M. V.; Vazquez-Corral, J.; Brown, M. S.; and Timofte, R. 2024. NILUT: Conditional Neural Implicit 3D Lookup Tables for Image Enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Ibrahim, R. I. 2014. DESIGN OF TWO-GLASS FOR APOCHROMATIC OPTICAL SYSTEM.
- Jiang, B.; Li, J.; Lu, Y.; Cai, Q.; Song, H.; and Lu, G. 2025. Efficient image denoising using deep learning: A brief survey. *Information Fusion*.
- Jiang, H.; Luo, A.; Liu, X.; Han, S.; and Liu, S. 2024. LightenDiffusion: Unsupervised Low-Light Image Enhancement with Latent-Retinex Diffusion Models. In *Proceedings of the European Conference on Computer Vision*.
- Jo, Y.; and Kim, S. J. 2021. Practical Single-Image Super-Resolution Using Look-Up Table. In *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Kang, S. B. 2007. Automatic Removal of Chromatic Aberration from a Single Image. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*.
- Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022. Denoising Diffusion Restoration Models. In *Advances in Neural Information Processing Systems*.
- Kim, B.-K.; and Park, R.-H. 2008. Automatic detection and correction of purple fringing using the gradient information and desaturation. In *2008 16th European Signal Processing Conference*.
- Kim, B.-K.; and Park, R.-H. 2010. Detection and correction of purple fringing using color desaturation in the xy chromaticity diagram and the gradient information. *Image Vision Comput.*
- Li, X.; Suo, J.; Zhang, W.; Yuan, X.; and Dai, Q. 2021. Universal and flexible optical aberration correction using deep-prior based deconvolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*.
- Li, X.; Zhang, B.; Sander, P. V.; and Liao, J. 2019. Blind Geometric Distortion Correction on Images Through Deep Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Van Gool, L.; and Timofte, R. 2021. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*.
- Lou, J.; Zhao, X.; Shi, K.; and Gu, S. 2025. Learning Pixel-adaptive Multi-layer Perceptrons for Real-time Image Enhancement. *arXiv preprint arXiv:2507.12135*.
- Malik, P.; and Karthik, K. 2020. Correction of complex purple fringing by green-channel compensation and local luminance adaptation. *IET Image Processing*.
- Maman, R.; Mualem, E.; Mazurski, N.; Engelberg, J.; and Levy, U. 2023. Achromatic imaging systems with flat lenses enabled by deep learning. *ACS Photonics*.
- Milster, T.; Kim, Y.; and Wang, Z. 2022. Type 2 longitudinal chromatic aberration from a high-harmonic MODE lens and color corrector. In *Proceedings of SPIE - The International Society for Optical Engineering*.
- Pont-Tuset, J.; Perazzi, F.; Caelles, S.; Arbeláez, P.; Sorkine-Hornung, A.; and Van Gool, L. 2017. The 2017 DAVIS Challenge on Video Object Segmentation. *arXiv:1704.00675*.
- Sanli, U.; Rodgers, G.; Zdora, M.-C.; Qi, P.; Garrevoet, J.; Falch, K.; Müller, B.; David, C.; and Vila-Comamala, J. 2023. Apochromatic X-ray focusing. *Light: Science & Applications*.
- Schuler, C. J.; Hirsch, M.; Harmeling, S.; and Schölkopf, B. 2015. Learning to deblur. *IEEE transactions on pattern analysis and machine intelligence*.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Tomaselli, V.; Guarnera, M.; Bruna, A.; and Curti, S. 2011. Automatic detection and correction of purple fringing artifacts through a window based approach. In *2011 IEEE International Conference on Consumer Electronics - Berlin*.
- Volkova, M. A.; Lutsiv, V. R.; Nedoshivina, L. S.; and Ivanova, A. A. 2019. Using the effect of longitudinal chromatic aberration for measuring distances from a single color photograph. *J. Opt. Technol.*
- Yang, C.; Jin, M.; Jia, X.; Xu, Y.; and Chen, Y. 2022a. Adaint: Learning adaptive intervals for 3d lookup tables on real-time image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Yang, C.; Jin, M.; Xu, Y.; Zhang, R.; Chen, Y.; and Liu, H. 2022b. Seplut: Separable image-adaptive lookup tables for real-time image enhancement. In *Proceedings of the European Conference on Computer Vision*. Springer.
- Yang, S.; Huang, B.; Zhang, Y.; Yu, D.; Yang, Y.; and Wong, N. 2025. DnLUT: Ultra-Efficient Color Image Denoising via Channel-Aware Lookup Tables. In *Proceedings of the Computer Vision and Pattern Recognition Conference*.
- Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Zeng, H.; Cai, J.; Li, L.; Cao, Z.; and Zhang, L. 2020. Learning image-adaptive 3d lookup tables for high performance photo enhancement in real-time. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.