

Nighttime Smartphone Reflective Flare Removal Using Optical Center Symmetry Prior

Yuekun Dai Yihang Luo Shangchen Zhou Chongyi Li* Chen Change Loy
 S-Lab, Nanyang Technological University

{YDAI005, c200211, s200094, chongyi.li, ccloy}@ntu.edu.sg

<https://ykdai.github.io/projects/BracketFlare>

Abstract

Reflective flare is a phenomenon that occurs when light reflects inside lenses, causing bright spots or a “ghosting effect” in photos, which can impact their quality. Eliminating reflective flare is highly desirable but challenging. Many existing methods rely on manually designed features to detect these bright spots, but they often fail to identify reflective flares created by various types of light and may even mistakenly remove the light sources in scenarios with multiple light sources. To address these challenges, we propose an optical center symmetry prior, which suggests that the reflective flare and light source are always symmetrical around the lens’s optical center. This prior helps to locate the reflective flare’s proposal region more accurately and can be applied to most smartphone cameras. Building on this prior, we create the first reflective flare removal dataset called BracketFlare, which contains diverse and realistic reflective flare patterns. We use continuous bracketing to capture the reflective flare pattern in the underexposed image and combine it with a normally exposed image to synthesize a pair of flare-corrupted and flare-free images. With the dataset, neural networks can be trained to remove the reflective flares effectively. Extensive experiments demonstrate the effectiveness of our method on both synthetic and real-world datasets.

1. Introduction

Lens flare artifacts can occur when a strong light enters the camera’s field of view, causing scattering and reflection within the lenses. Unlike radial-shaped scattering flares, reflective flares appear as groups of bright spots [5, 7] or a “ghosting effect” like the examples in Fig. 1. With multiple reflections, the light source creates multiple projections at various positions behind the lens. If the projection is far from the focal plane, the defocused effect can brighten the whole sensor and degrade the photo’s quality as shown in Fig. 2. In near-focus situations, images of the light source

*Corresponding author.

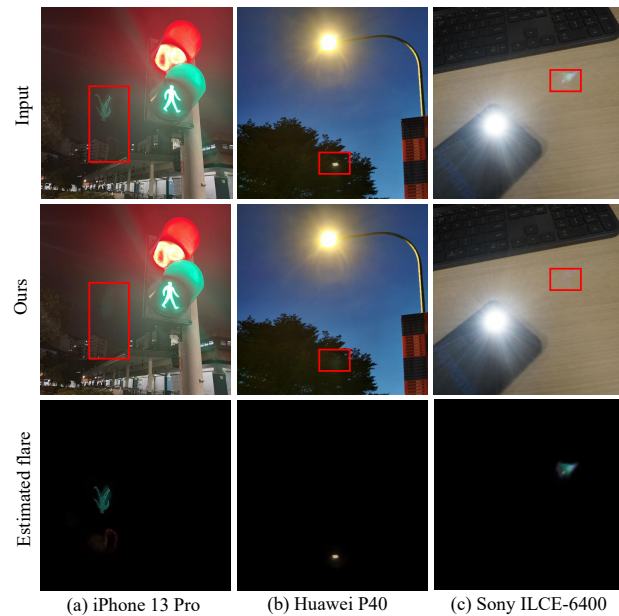


Figure 1. Reflective flare removal with our proposed dataset and prior. The network trained on our dataset can effectively separate the reflective flare and flare-free images from a flare-corrupted image. Our method can generalize well to different types of smartphones and even cameras with specific types of lenses.

can result in a line of unfocused, bright blobs on the captured photo.

Applying anti-reflective (AR) coating is a common technique for reducing internal reflections [22]. The coating is effective for unfocused spots with low luminance, but cannot mitigate a bright spot with high luminance at the focal plane. Even with a well-designed lens system, a bright reflective spot may remain in the photo, known as the ghosting effect [10]. Reflective flare is common on all types of smartphone lenses, and certain types of smartphones (e.g., the iPhone series) can produce a large area of inverted reflective flare from neon lights or large, luminous standing signs, as shown in Fig. 1. The phenomenon becomes more pronounced when capturing a photo at night with multiple lights in the scene. Numerous bright spots and inverted re-

fections manifest and propagate in different directions with the lens orientation. Therefore, reflective flare removal algorithms are highly desirable.

Removing reflective flare is a difficult task due to the vast diversity of reflective flares that can be caused by different types of light sources and lenses. Reflective flares often appear as bright spots that can resemble streetlights from a distance, making it challenging to differentiate between reflective flares and other light sources. Current methods for removing flare spots mainly rely on feature detection techniques like Speeded-Up Robust Features [2] (SURF) or Scale-Invariant Feature Transform [12] (SIFT) to extract spot regions [1, 3, 19]. These methods, however, struggle to effectively detect blobs on complex backgrounds and only work for specific spot patterns. Moreover, these methods are ineffective in removing reflective flares caused by lamp tubes and LED matrices. Many learning-based methods and datasets for reflective flare removal have been proposed [5, 15, 18, 22]. Although these datasets contain subsets of reflective flares, they are not sufficient for good generalization. For instance, the reflective flare subset from Wu *et al.* [22] is captured using the same lens and light source, making it difficult to generalize to different flares in real-world scenarios. Dai *et al.* [5] propose a synthetic dataset with 10 types of reflective flares, but training on this dataset tends to mislead a network to erroneously remove other tiny bright spots like street lights and bright windows. As a result, it remains challenging for existing algorithms to accurately detect and remove various types of reflective flares while preserving details in other regions.

To overcome the challenges of detecting and removing reflective flares, we propose a novel optical center symmetry prior, which suggests that the main spot of the reflective flare and the light source are always symmetrical around the lens’s optical center in captured images. This prior applies to most smartphone cameras and some professional cameras. With this prior, we can easily detect the possible locations of reflective flares. We also found that the flare spots always have the same patterns as the brightest regions of the light source. Therefore, we can synthesize reflective flares by using the light source’s patterns at low exposure. Based on these characteristics of reflective flares, we develop the first reflective flare removal dataset, named *BracketFlare*. Specifically, we use continuous bracketing to capture the light source’s pattern on an underexposed image and select the normally exposed image to create the light source and the natural image. Next, we rotate the light source pattern 180 degrees around the optical center, apply blur and color augmentation, and use it to create synthetic reflective flares. We can add the reflective flare to the normally exposed image to produce paired flare-corrupted and flare-free images.

Thanks to the optical center symmetry prior, we further propose an end-to-end pipeline for reflective flare re-

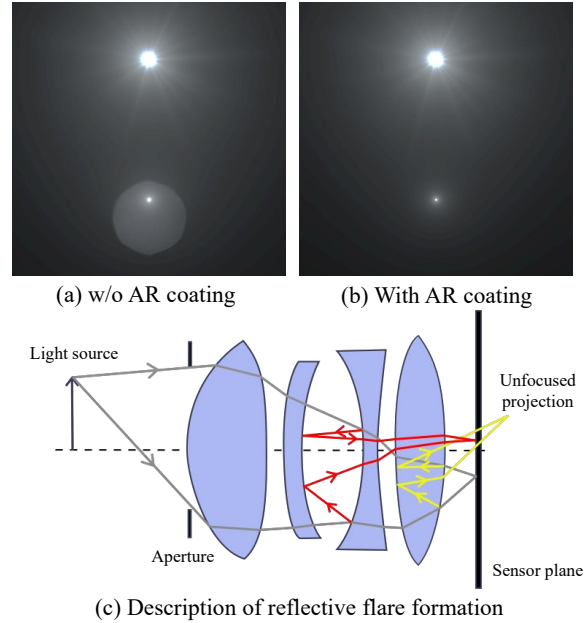


Figure 2. Schematic diagram of reflective flare formation. Reflective flares are generated by multiple reflections inside the lens. In (c), yellow lines show the generation principle of the aperture-shaped reflective flare. Red lines represent the formation of the bright spot. Since the unfocused reflective flare (yellow line, e.g., (a)) can be effectively alleviated by AR coating. In this work, we mainly focus on removing the focused reflective flare (red line, e.g., (b)) from night images.

moval, which can easily apply standard neural networks designed for image restoration. Specifically, to demonstrate the effectiveness of our dataset and pipeline, we train two CNN-based networks [4, 24] and two Transformer-based networks [21, 23] as baselines. These models trained on our dataset effectively suppress the reflective flare and generalize well to real-world scenarios. The contributions of our work are as follows:

- We propose a new optical center symmetry prior that uncovers crucial properties of reflective flares on smartphones’ cameras, including their locations of occurrence and patterns.
- Based on this prior, we explore a new method for constructing a reflective flare removal dataset called *BracketFlare*. This dataset includes various types of light sources and diverse indoor and outdoor scenes, providing a solid foundation for this task.
- We design an end-to-end reflective flare removal pipeline that incorporates the optical center symmetry prior to both training data generation and network structures. Extensive experiments demonstrate the essential role of our dataset and prior in reflective flare removal, which effectively generalizes to different night scenes

2. Related Work

Scattering flare removal. Most flare removal methods mainly focus on removing the scattering flares and glare effects. Based on the glow smooth prior, unsupervised methods [8, 16] are proposed to separate the glare effect from the flare-corrupted images. For supervised flare removal, Wu *et al.* [22] and Dai *et al.* [5] have respectively proposed two flare image datasets. These flare images can be added to flare-free images to synthesize the flare-corrupted and flare-free image pairs, making training an end-to-end flare removal network possible. Due to the difficulty in collecting real paired data, Qiao *et al.* [15] designed a pipeline to train the flare removal network with unpaired data following the idea of Cycle-GAN [27]. Besides, artifacts from under-display cameras [6] and nighttime haze [11] share similar attributes with the scattering flares. Thus, the related methods can also alleviate scattering flares artifacts.

Reflective flare removal. Traditional reflective flare removal is typically accomplished through feature detection and image inpainting. Flare detection in these approaches [1, 3, 13, 19] assumes that reflective flares are bright spots with specific shapes. Such an assumption is fragile given nighttime situations with diverse light sources. Since reflective flare removal can be viewed as a specific kind of reflection removal [26], one can adopt learning-based methods. However, existing datasets for this task, such as those proposed by Wu *et al.* [22], Dai *et al.* [5], and Qiao *et al.* [15], are not specifically designed for reflective flare removal and lack diversity in terms of light sources. These limitations hinder deep learning-based methods from effectively addressing real-world reflective flares, and may even lead to the removal of other tiny bright spots such as street lights.

3. Optical Center Symmetry Prior

3.1. Definition for the prior

When pointing the camera at a light source, strong light can sometimes create lens flares in photos. In well-designed lens systems with AR coatings, reflective flares appear as bright spots on the photo, rather than generating a light fog. Through our investigation of most smartphones' main cameras, we discover that the bright spots of reflective flares and light sources are always symmetrical around the lens's optical center, as shown in Fig. 3. By aligning the light source and optical center, we can easily locate the reflective flares, eliminating the need to search the entire image for possible occurrences. Besides, since the main bright spots are caused by focused reflective flares, they share the same attributes as the light source. For example, the reflective flares may display clear details for each small LED lamp bead, as shown in Fig. 3(b) and Fig. 3(c).




Figure 3. Strong light source may project a bright reflective flare on the photo. The reflective flares and corresponding light sources are symmetric around the optical center for most smartphone cameras. The yellow cross represents the optical center.

3.2. Physical formulation

To illustrate the formation of optical center symmetry prior, we suppose that an optical ray \mathbf{r} is emitted from the light source, as shown in Fig. 4. This ray can be described as $\mathbf{r} = [h \ \theta]^\top$. The states of this ray after each lens surface can be calculated by multiplying transfer matrices [10, 14]. We define the translation matrix at layer i as \mathbf{T}_i , and the refraction matrix at spherical surface with radius R_i and refractive index n_i, n_{i+1} as \mathbf{F}_i . For the reflection at the i th layer with radius R_i , the reflection matrix can be stated as \mathbf{R}_i . These matrices can be stated as:

$$\mathbf{T}_i = \begin{pmatrix} 1 & d_i \\ 0 & 1 \end{pmatrix}, \mathbf{R}_i = \begin{pmatrix} 1 & 0 \\ \frac{2}{R_i} & 0 \end{pmatrix}, \mathbf{F}_i = \begin{pmatrix} 1 & 0 \\ \frac{n_{i+1} - n_i}{n_{i+1} R_i} & \frac{n_i}{n_{i+1}} \end{pmatrix} \quad (1)$$

In Fig. 4, h_1 represents the position of a real image that focuses on the sensor through the lens' internal reflection. We define the ray's forward propagation matrix as \mathbf{A} and \mathbf{C} and the reflection process as \mathbf{B} . They are defined as:

$$\begin{aligned} \mathbf{A} &= \mathbf{T}_4 \mathbf{F}_3 \mathbf{T}_3 \mathbf{F}_2 \mathbf{T}_2 \mathbf{F}_1 \mathbf{T}_1 \mathbf{F}_0 \mathbf{T}_0 \\ \mathbf{B} &= \mathbf{T}_4 \mathbf{F}_3 \mathbf{T}_3 \mathbf{R}_2^{-1} \mathbf{T}_3 \mathbf{F}_3^{-1} \mathbf{T}_4 \mathbf{R}_4 \\ \mathbf{C} &= \mathbf{T}_8 \mathbf{F}_8 \mathbf{T}_7 \mathbf{F}_7 \mathbf{T}_6 \mathbf{F}_5 \mathbf{T}_5 \mathbf{F}_4 \end{aligned} \quad (2)$$

With these matrices, the ray after reflection in the focal plane \mathbf{r}_1 and direct imaging \mathbf{r}_0 can be written as:

$$\begin{pmatrix} h_1 \\ \theta_1 \end{pmatrix} = \mathbf{CBA} \begin{pmatrix} h \\ \theta \end{pmatrix}, \begin{pmatrix} h_0 \\ \theta_0 \end{pmatrix} = \mathbf{CA} \begin{pmatrix} h \\ \theta \end{pmatrix} \quad (3)$$

For the focused image, h_0 and h_1 will not vary with the change of the angle of incidence θ . Thus, we can deduce from Eq. (3) that both h_0 and h_1 are linear with h . This

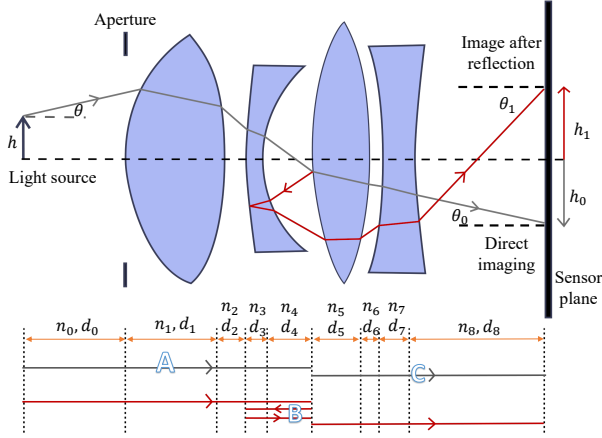


Figure 4. Physical explanation of reflective flare formation. In this figure, h and θ indicate the height and the angle of incidence of the light source. d and n denote the thickness and refractive index of each layer. A , B , and C represent the transfer matrices for each part of the lens system. The red line represents the imaging after reflection and the grey line shows the path of direct imaging. This figure depicts the situation with reflection in the 3rd and 5th air-lens layer, in real-world situations, the reflections can happen at any place including the lens protector and CMOS layer.

means that the line between the light source and the reflective flare will pass through the optical center and the ratio between the h_0 and h_1 is a constant. We define this ratio as k with $h_1 = kh_0$. Thus, we can calculate the position of the reflective flare, if this constant k is known. Due to the lens design for most smartphone cameras, the value of k is always -1. It leads to the optical center symmetry prior stated above.

3.3. Experimental statistics

To demonstrate the effectiveness of our proposed prior in real-world situations, we conducted experiments using the main cameras of three smartphones: the Huawei P40, iPhone 13 Pro, and ZTE Axon 20 5G. In dark scenes, we captured 20 photos using each camera and marked the positions of the light source and reflective flare relative to the center of the image. As shown in Fig. 5, we observed that for all three cameras, the position of the reflective flare is linearly related to the position of the light source, with a slope of 1. However, since the light source may often create overexposed regions, manually marking the position of the light source can still introduce some errors. Considering this effect, we conclude that the results from all three phones align with our proposed prior.

4. Data Collection

Although the optical center prior can guide us in finding the reflective flare, the lack of paired images still presents a challenge for training an effective model for flare re-

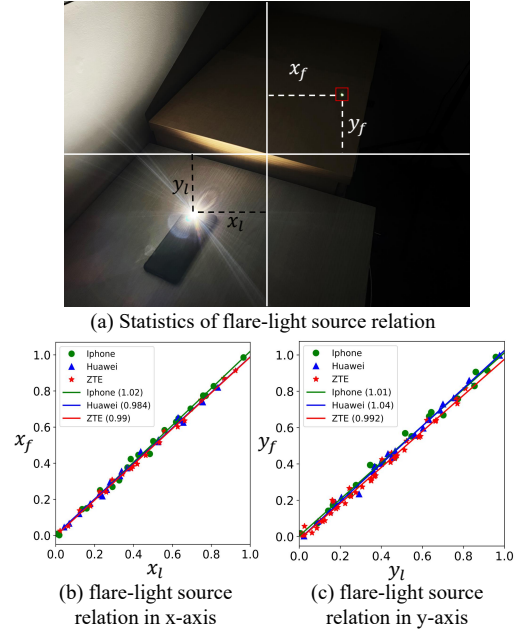


Figure 5. Experimental statistics of the optical center symmetry prior. We capture images with reflective flares in a dark environment and mark the position of the light source and reflective flares. x_f and x_l represent the x-coordinate of flare and light source and so as y_f and y_l . The origin is set to be the center of the image. The relationship between the flare and the light source’s absolute values of coordinates are plotted above. From figures (b) and (c), the coordinate relationships are linear and the slope is 1, which confirms our prior.

moval. As reflective flares are caused by the lens system, it is physically impossible to capture real flare-free and flare-corrupted image pairs. Moreover, given the huge variety of nighttime lights, it would be tedious to synthesize reflective flares for each type of light source.

To address the challenge of generating paired images for reflective flare removal, we propose a semi-synthetic dataset, named *BracketFlare*. Since reflective flares always have the same pattern as the brightest region of the light source, we reduce the exposure time to capture the patterns of the light source’s brightest region, which are used to synthesize the reflective flares. Specifically, we mount a professional camera (Sony A7 III) on a tripod and use continuous bracketing to capture five photos of the same static scene with a step of 3 EV as a group. We then select the normally exposed image and the lowest exposure image from each group to compose the pair. To increase the diversity and realism of our dataset, we choose scenes with different light sources that are bright enough to leave reflective flares. Finally, we collect a dataset consisting of 440 pairs of 4K images as training data and 40 images as test data. A general overview of our dataset is depicted in Fig. 6, showing that the dataset contains diverse and commonly seen flare patterns.

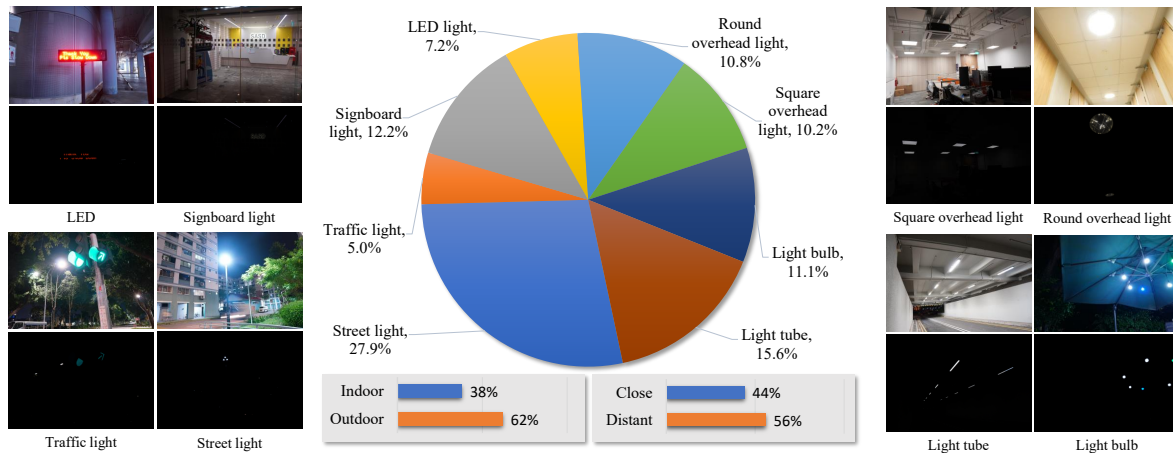


Figure 6. Visualization and distribution of our dataset. It contains different types of light sources in diverse scenes. Based on light source types, we classify them into eight categories with different flare patterns.

5. Proposed Method

5.1. Postprocessing our dataset

To increase the diversity of our dataset, we generate paired flare-corrupted and flare-free images on the fly with augmentation. As depicted in Fig. 7(a), we first resize the 4K image to 800×1200 . To simulate the possible misalignment of the image’s center and optical center, we introduce a slight translation with $t \sim U(0, 15)$ pixels to the under-exposed light source image. Then, we randomly crop the paired light source image and the normally exposed image with a size of 512×512 . Since our captured images are already gamma-encoded, we follow the settings of Wu *et al.* [22] and linearize them by applying a gamma correction, i.e., $I_g = I^\gamma$ with γ randomly sampled from $[1.8, 2.2]$. All our operations are performed in the linear space.

Synthetic reflective flare. As lenses reflect light of different wavelengths with varying reflectivity, reflective flares often appear in different colors. To replicate this, we use a random color factor to multiply the low-exposed image. To simulate an unfocused situation, we also add random blur. In accordance with the optical center symmetry prior, we rotate the light source image to synthesize the reflective flare. Finally, we obtain the synthetic reflective flare by multiplying the flare image with an opacity.

Flare-free image. The ISO is fixed for all images in our continuous bracketing. Therefore, our normally exposed images always share a long exposure time and have little noise. To simulate real night scenes, we add Gaussian noise to the images. Finally, we randomly apply a small color offset to the flare-free image. These operations aid in simulating light fog and glare effects in real-world scenarios.

Flare-corrupted image. To create flare-corrupted images, we add the synthetic reflective flare to the flare-free images. Then, we apply gamma correction to the images with the

gamma value obtained by taking the reciprocal of the previously sampled gamma correction value γ . The resulting gamma-encoded flare-corrupted image, flare-free image, and flare image are used to train the neural networks.

5.2. Flare removal pipeline

As the optical center symmetry prior is a global prior, it is difficult for CNN-based networks to learn it, and even Transformer-based networks struggle to learn such symmetry. Therefore, it is crucial to incorporate this prior into the network with specialized designs.

To improve the network’s ability to learn the optical center symmetry prior, as shown in Fig. 7(b), we apply a gamma correction of $I_g = I^\gamma$ to the input flare-corrupted image with $\gamma = 10$. This operation extracts the nearly saturated areas of the light source in each channel. We then rotate this nearly saturated image by 180 degrees to obtain a prior image, which is concatenated with the flare-corrupted image to form a six-channel input. This rotated prior image provides an initial approximation of the reflective flare’s pattern, which we expect the network to refine and enhance. The network’s output is set to six channels, with three channels for the estimated flare-free image and three channels for the estimated flare. Our synthetic data provides supervision for all six channels, which facilitates network training.

5.3. Loss function

With the supervision of flare image I_F and flare-free image I_B , our final loss is the combination of background loss \mathcal{L}_{bg} , flare loss \mathcal{L}_{flare} , and reconstruction loss \mathcal{L}_{rec} . It can be written as:

$$\mathcal{L} = \mathcal{L}_{rec} + \mathcal{L}_{flare} + \mathcal{L}_{bg}. \quad (4)$$

We define our flare removal network as f_θ . It takes flare-corrupted image I_{input} and its prior image I_{prior} as inputs.

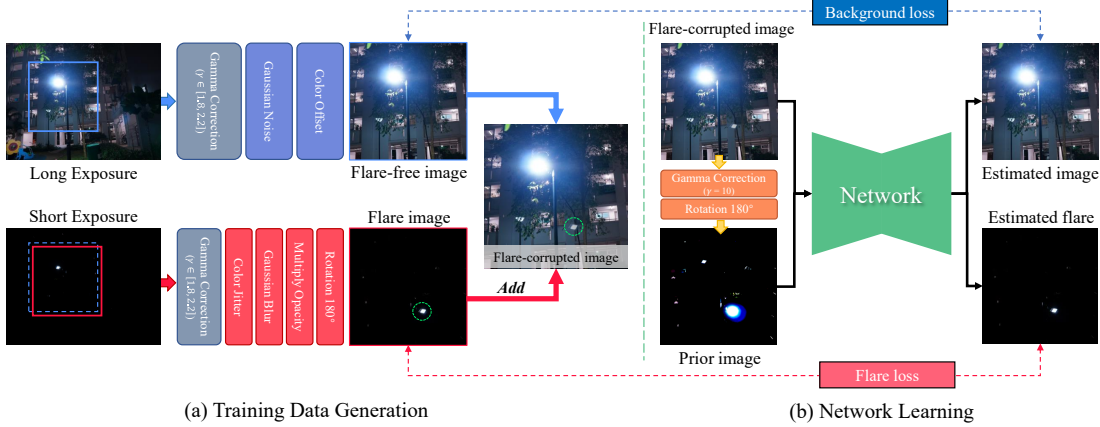


Figure 7. The pipeline of our data processing module and flare removal network. The flare image and a flare-free image from our paired dataset are resized and randomly cropped. After data augmentation, they are added to synthesize the corresponding flare-corrupted image. To encode the optical symmetry prior, the input flare-corrupted image is center flipped as the prior image. We also apply gamma correction to extract the light source region. Then, the prior image is concatenated with the flare-corrupted image as the input with 6 channels. Finally, the network predicts an output with 6 channels (a flare-free image with 3 channels and an estimated flare with 3 channels).

The flare-free image and flare can be estimated by:

$$\hat{I}_B, \hat{I}_F = f_{\theta}(I_{input}, I_{prior}). \quad (5)$$

With the estimated flare and flare-free images, the L_1 reconstruction loss can be expressed as:

$$\mathcal{L}_{rec} = |I_{input} - \hat{I}_F \oplus \hat{I}_B|, \quad (6)$$

where \oplus denotes an element-wise addition operation in the linearized gamma-decoded domain.

For the background loss, it consists of a default \mathcal{L}_1 , a region-aware masked \mathcal{L}_{mask} term, and a perceptual term \mathcal{L}_{vgg} with a pre-trained VGG-19 network [17]. Compared with previous works [5, 22, 26], the main difference is the \mathcal{L}_{mask} . Since the regions of reflective flares are always small, a simple L_1 loss may ignore these regions, leading to the local optimum. To encourage the network to focus on restoring the flare-corrupted regions, we calculate the flare’s mask $I_M (> 0.2)$ from the ground truth of flare image I_F . This masked loss can be written as:

$$\mathcal{L}_{mask} = |I_M * (\hat{I}_B - I_B)|. \quad (7)$$

Overall, the background loss \mathcal{L}_{bg} is defined as:

$$\mathcal{L}_{bg} = w_1 \mathcal{L}_1 + w_2 \mathcal{L}_{vgg} + w_3 \mathcal{L}_{mask}, \quad (8)$$

where w_1, w_2, w_3 are respectively set to 0.5, 0.1, and 20.0 in our experiments. Flare loss \mathcal{L}_{flare} has the identical form as the background loss.

6. Experiments

Evaluation on synthetic data. To evaluate our method, we collect 40 paired data using the process mentioned in

Table 1. Percentage of users favoring our results (our dataset + MPRNet[†]) vs. previous work. Our method outperforms existing methods on both indoor and outdoor images.

Method	Outdoor	Indoor
Reflection removal [26]	3.0%	8.5%
Flare spot removal [1]	4.5%	0.0%
Flare removal (Wu et al.) [22]	0.5%	3.5%
Flare removal (Dai et al.) [5]	1.5%	2.5%
Ours	90.5%	85.5%

Sec. 4 to generate a synthetic test dataset, in which ground truth images are available. We adopt PSNR, SSIM [20], and LPIPS [25] as our evaluation metrics.

Evaluation on real data. To show our method’s generalization ability, we also collect a dataset with the photos captured by different types of mobile phones, including iPhone 12 pro, iPhone 13 pro, iPhone 14 pro, Huawei mate 30, Huawei nova 5i Pro, Huawei P40, Samsung Galaxy s22, Samsung Galaxy Note9, ZTE Axon 20 5G and Oppo Reno4 SE 5G. For real cases, we cannot obtain the ground truth. Since reflective flares always locate in small regions, they are hard to be evaluated by existing non-reference image quality assessment. These metrics emphasize more on global image quality and are not specially designed for our task. Therefore, we provide more visual results and conduct a user study to evaluate the performance.

6.1. Comparison with different networks

To validate the performance of our pipeline, we train different CNN-based and Transformer-based networks with our data. We use the default network settings of MPRNet [24], HINet [4], and Uformer [21] and reduce the feature channels of Restormer [23] from 48 to 24 due to GPU memory constraint. In our experiments, these models are



Figure 8. Visual comparison of different flare removal and reflection removal methods on real-world nighttime flare images.




Figure 9. Visual comparison of different flare removal and reflection removal methods on synthetic nighttime flare images.

trained on an NVIDIA Geforce RTX3090 with a batch size of 2. We use Adam optimizer [9] with a learning rate of $1e-4$ and train a model for 500k iterations. Besides, we adopt a MultiStepLR with a milestone of 200k and set the gamma to 0.5. These models are evaluated using the synthetic test dataset as shown in Table 4. From experiments, we find that MRPNet[†] [24] performs better than other networks. Thus, we treat MRPNet[†] [24] as our baseline method and conduct the following ablation study.

6.2. Comparison with previous works

We compare our method with different flare removal datasets [5, 22] and reflection removal dataset [26] on both real and synthetic data. We use the pre-trained models of Dai *et al.* [5] and Zhang *et al.* [26]. Since Wu *et al.* [22] do not release their model, we retrain their network using their provided data and code. Besides, we also compare our method with the traditional flare spot removal method [1]. As shown in Table 2, since most of these methods are not specially designed for reflective flare removal, they cannot

Table 2. Quantitative comparison on synthetic test data.

Method	PSNR	SSIM	LPIPS	Masked PSNR
Input	37.30	0.990	0.025	21.68
Reflection removal [26]	22.02	0.830	0.074	22.88
Flare spot removal [1]	38.29	0.990	0.024	22.53
Flare removal (Wu <i>et al.</i>) [22]	26.13	0.895	0.055	21.89
Flare removal (Dai <i>et al.</i>) [5]	27.49	0.942	0.050	21.85
Ours	48.41	0.994	0.004	32.09

accurately remove the reflective flare while retaining image details simultaneously. Although Asha *et al.* [1] can detect the bright spots, as shown in Fig. 8, it may also remove the light source which is quite common at night.

User Study. To evaluate these methods in real-world situations, we also conduct a user study with 20 participants with 20 images randomly selected from our real data in Table 1.

6.3. Ablation study

Loss function. We conduct an ablation study to prove the necessity of each loss function used in our baseline. In our experiments, we train the model without adding \mathcal{L}_{flare} ,



Figure 10. Visual comparison of flare removal on real-world nighttime flare images that are captured by different smartphones.

Table 3. Ablation study for different loss functions.

Loss	PSNR	SSIM	LPIPS	Masked PSNR
w/o \mathcal{L}_{flare}	47.36	0.993	0.004	30.63
w/o \mathcal{L}_{rec}	48.18	0.994	0.003	31.33
w/o \mathcal{L}_{vgg}	48.05	0.994	0.004	31.39
w/o \mathcal{L}_{mask}	48.01	0.993	0.003	31.03
Full	48.41	0.994	0.004	32.09




Figure 11. Visual comparison with the network’s performance with and w/o optical center symmetry prior.

\mathcal{L}_{rec} , \mathcal{L}_{vgg} or \mathcal{L}_{mask} . We report the results in Table 3. The ablation study of \mathcal{L}_{flare} proves the significance of bringing supervision for both flare and background images. Besides, when the \mathcal{L}_{mask} is removed, it degrades the reconstruction results in the flare-corrupted regions, which verifies the importance of focusing on flare regions.

Prior. To illustrate the performance of introducing the optical symmetry prior by concatenating a prior image, we train these baseline restoration networks with the original input. Since HINet [4] and MPRNet [24] cannot support input and output with different channels, we concatenate a black image with no information and keep other settings the same. The quantitative results are shown in Table 4, and visual comparisons on real-world images are presented in Fig.

Table 4. Quantitative comparison for different network structures on synthetic test data. ‘†’ indicates that the network is equipped with the proposed prior.

Network	PSNR	SSIM	LPIPS	Masked PSNR
MPRNet [24]	43.60	0.992	0.011	27.64
Uformer [21]	42.57	0.989	0.009	27.77
HINet [4]	43.78	0.991	0.012	27.73
Restormer [23]	42.76	0.991	0.012	27.27
MPRNet†	48.41	0.994	0.004	32.09
Uformer†	47.47	0.991	0.003	31.57
HINet†	48.03	0.994	0.003	30.88
Restormer†	48.11	0.994	0.004	31.07

11. It shows that embedding this prior in the network can boost the performance of quantitative results considerably. In Fig. 11, we observe that this prior helps the network to locate and remove the reflective flare precisely.

7. Conclusion

We have presented a novel approach to tackle the challenge of removing reflective flares in nighttime photography. Our approach is inspired by the optical center symmetry prior, which suggests that reflective flares and their corresponding light sources are always symmetrical around the optical center of the lens system in smartphone cameras. Based on the observation that the patterns of reflective flares are always similar to the brightest region of the light source, we introduce a dataset named *BracketFlare*. We further propose a reflective flare removal pipeline that encodes the optical center symmetry prior into the network. Experimental results demonstrate that our proposed methods are effective in removing reflective flares in real-world scenarios and have the potential for wide-ranging applications.

Acknowledgement: This study is supported under the RIE2020 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) Funding Initiative, as well as cash and in-kind contribution from the industry partner(s).

References

- [1] C S Asha, Sooraj Kumar Bhat, Deepa Nayak, and Chaithra Bhat. Auto removal of bright spot from images captured against flashing light source. In *ICDC*, 2019. 2, 3, 6, 7
- [2] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer Vision and Image Understanding*, 110(3):346–359, 2008. 2
- [3] Floris Chabert. Automated lens flare removal, 2014. Technical Report, Department of Electrical Engineering, Stanford University. 2, 3
- [4] Liangyu Chen, Xin Lu, Jie Zhang, Xiaojie Chu, and Chengpeng Chen. Hinet: Half instance normalization network for image restoration. In *CVPRW*, 2021. 2, 6, 8
- [5] Yuekun Dai, Chongyi Li, Shangchen Zhou, Ruicheng Feng, and Chen Change Loy. Flare7k: A phenomenological nighttime flare removal dataset. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 1, 2, 3, 6, 7
- [6] 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 network. In *CVPR*, 2021. 3
- [7] Matthias Hullin, Elmar Eisemann, Hans-Peter Seidel, and Sungkil Lee. Physically-based real-time lens flare rendering. In *ACM SIGGRAPH*, 2011. 1
- [8] Yeying Jin, Wenhao Yang, and Robby T Tan. Unsupervised night image enhancement: When layer decomposition meets light-effects suppression. In *ECCV*. Springer, 2022. 3
- [9] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 7
- [10] Sungkil Lee and Elmar Eisemann. Practical real-time lens flare rendering. *Computer Graphics Forum*, 32(4):1–6, 2013. 1, 3
- [11] Yu Li, Robby T. Tan, and Michael S. Brown. Nighttime haze removal with glow and multiple light colors. In *ICCV*, 2015. 3
- [12] David G Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 2
- [13] Andreas Nussberger, Helmut Grabner, and Luc Van Gool. Robust aerial object tracking in images with lens flare. In *ICRA*, 2015. 3
- [14] Frank L Pedrotti, Leno M Pedrotti, and Leno S Pedrotti. *Introduction to Optics*. Cambridge University Press, 2017. 3
- [15] Xiaotian Qiao, Gerhard P. Hancke, and Rynson W. H. Lau. Light source guided single-image flare removal from unpaired data. In *ICCV*, 2021. 2, 3
- [16] Aashish Sharma and Robby T. Tan. Nighttime visibility enhancement by increasing the dynamic range and suppression of light effects. In *CVPR*, 2021. 3
- [17] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2014. 6
- [18] Qilin Sun, Ethan Tseng, Qiang Fu, Wolfgang Heidrich, and Felix Heide. Learning rank-1 diffractive optics for single-shot high dynamic range imaging. In *CVPR*, 2020. 2
- [19] Patricia Vitoria and Coloma Ballester. Automatic flare spot artifact detection and removal in photographs. *Journal of Mathematical Imaging and Vision*, 61:515–533, 2019. 2, 3
- [20] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 6
- [21] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *CVPR*, 2022. 2, 6, 8
- [22] 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. 1, 2, 3, 5, 6, 7
- [23] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 2022. 2, 6, 8
- [24] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, 2021. 2, 6, 7, 8
- [25] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. 6
- [26] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *CVPR*, 2018. 3, 6, 7
- [27] Junyan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, 2017. 3