Flare7K: A Phenomenological Nighttime Flare Removal Dataset (Supplementary Material)

Yuekun Dai Chongyi Li Shangchen Zhou Ruicheng Feng Chen Change Loy
S-Lab, Nanyang Technological University
{YDAI005, chongyi.li, s200094, ruicheng002, ccloy}@ntu.edu.sg
https://Nukaliad.github.io/projects/Flare7K

In this supplementary material, we present additional details of the proposed Flare7K dataset and experimental settings and show more results.

- Physics of different kinds of scattering flares
- Experimental details
- More visual results
- Extension application details
- Social impacts
- Dataset documentation

1 Physics of different kinds of scattering flares

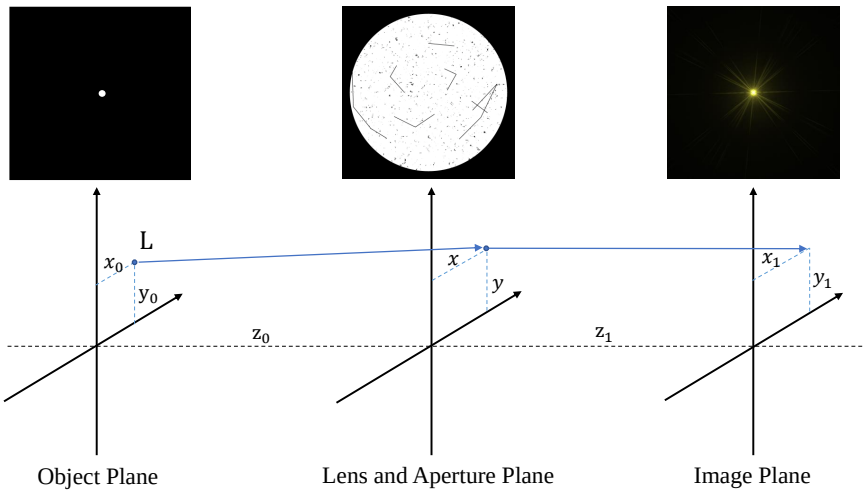

Figure 1: Illustration of a simplified lens system. In this figure, L represents a point light source. In the lens and aperture plane, the light passes through the dirty aperture and lens system, leaving a scattering flare on the image plane.

In this section, we use a simplified Fourier optics model to illustrate how different kinds of scattering flares occur. A basic lens system can be viewed as a combination of one convex lens, one aperture, and an image plane as shown in Figure 1. We set the optical center as the origin of a coordinate system.

36th Conference on Neural Information Processing Systems (NeurIPS 2022) Track on Datasets and Benchmarks.

Then, the light source's position is $(x_0, y_0, -z_0)$. In Fourier optics, a complex-valued divergent spherical wave from light source L can be written as:

$$\widetilde{U}(x,y) = \frac{A_0}{\sqrt{(x-x_0)^2+(y-y_0)^2+z_0^2}} \exp\left(ik\sqrt{(x-x_0)^2+(y-y_0)^2+z_0^2}\right), \tag{1}$$

where A_0 represents the amplitude of the light source, $\widetilde{U}(x,y)$ is the complex-valued wave function at the aperture plane, and k is the wavenumber corresponding to a wavelength λ with $k = 2\pi/\lambda$. We suppose that this problem satisfies paraxial approximation:

$$\rho_0 = \sqrt{x_0^2+y_0^2}, \;\; \rho = \sqrt{x^2+y^2}, \;\; \rho_0, \rho << z_0. \tag{2}$$

Under paraxial approximation, the divergent spherical wave can be written as:

$$\widetilde{U}(x,y) = \frac{A_0 e^{ikr_0}}{z_0} \exp\left(ik\frac{x^2+y^2}{2z_0}\right) \exp\left(-ik\frac{xx_0+yy_0}{z_0}\right), \tag{3}$$

$$r_0 = \sqrt{x_0^2+y_0^2+z_0^2} \approx z_0 + \frac{x_0^2+y_0^2}{2z_0}. \tag{4}$$

In an aperture plane, the transformation function of the wavefront can be represented as $\widetilde{T}(x,y)$. It is a combination of aperture function $\widetilde{A}_\lambda(x,y)$ and a lens function $\widetilde{T}_L(x,y)$. Supposing the focus of the lens is f and the lens is ideal. Complex-valued lens function satisfies:

$$\widetilde{T}_L(x,y) = \exp\left(-ik\frac{x^2+y^2}{2f}\right). \tag{5}$$

For a real-world aperture with dirt or scratches, the aperture function can be represented as the multiplication of clear aperture function $A_a(x,y)$ and different types of dirt function $\widetilde{A}_d(x,y)$:

$$\widetilde{A}_\lambda(x,y) = A_a(x,y) \prod \widetilde{A}_d(x,y), \tag{6}$$

$$A_a(x,y) = \begin{cases} 1 & \text{if } x^2+y^2 < r^2 \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

After refracting in the lenses, the amplitude and phase for each position at output plane of aperture can be characterized by pupil function $\widetilde{P}_\lambda(x,y)$:

$$\widetilde{P}_\lambda(x,y) = \widetilde{T}_L(x,y)\widetilde{A}_\lambda(x,y)\widetilde{U}(x,y). \tag{8}$$

From the equations above, the pupil function can also be expressed as:

$$\widetilde{P}_\lambda(x,y) = \frac{A_0\widetilde{A}_\lambda(x,y)e^{ikr_0}}{z_0} \exp\left(ik(x^2+y^2)\left(\frac{1}{2z_0}-\frac{1}{2f}\right)\right) \exp\left(-ik\frac{xx_0+yy_0}{z_0}\right). \tag{9}$$

On the image plane, the final image plane function can be represented as integral of the pupil function. For a point (x_1, y_1, z_1) located at image plane, if image plane is in focus with $z_0^{-1} + z_1^{-1} = f^{-1}$, the optical field can be written as:

$$\widetilde{U}_I(x_1,y_1) = \iint \widetilde{P}_\lambda(x,y) \exp\left(ik\sqrt{(x_1-x)^2+(y_1-y)^2+z_1^2}\right)dxdy \tag{10}$$

$$= \widetilde{A}_c \iint \widetilde{A}_\lambda(x,y) exp\left(-ik\frac{xx_0+yy_0}{z_0}\right) \exp\left(-ik\frac{xx_1+yy_1}{z_1}\right)dxdy, \tag{11}$$

where $\widetilde{A}_c$ is a constant coefficient. After adjusting the origin of x_1 and x_2, Equation (11) can be viewed as a standard Fourier transformation. Thus, the point spread function (PSF) which is the square of the amplitude of the image plane's optical field can be written as:

$$PSF_\lambda = |\mathcal{F}\{\widetilde{A}_\lambda(x,y)\}|^2. \tag{12}$$

The $\mathcal{F}$ here represents Fourier transformation. Since stains with depth may bring phase shift for the aperture function, the PSF_λ may vary with the wavelength λ of the light source. Also, the lens is not ideal and may have dispersion. To make it simple, we just suppose that all the scratches and dirt are

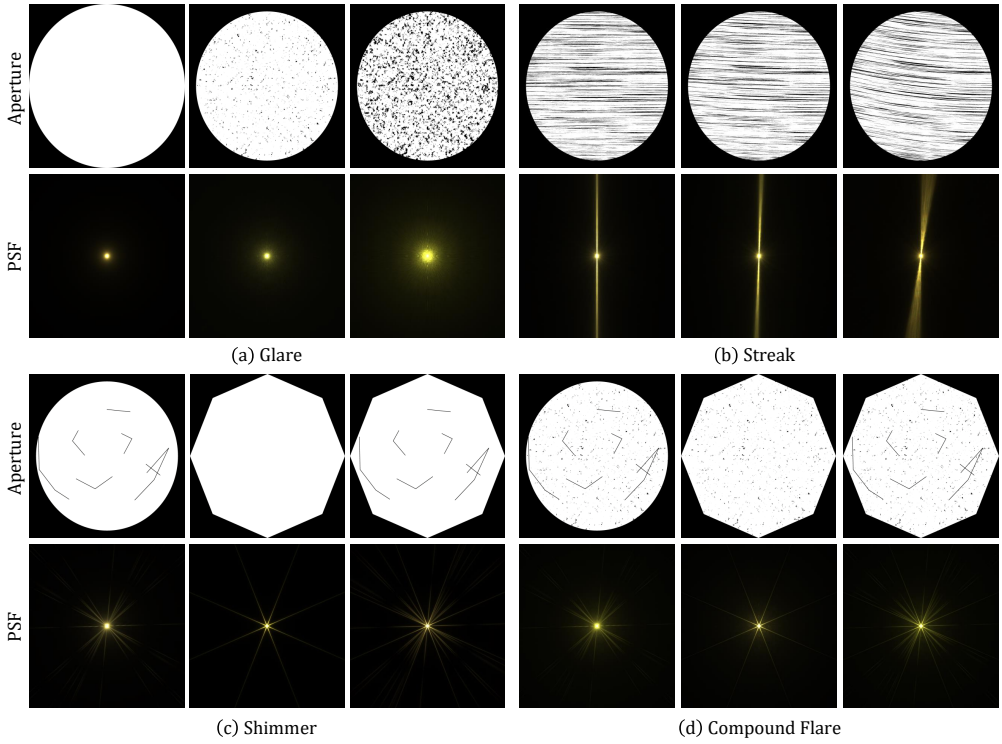

Figure 2: Different types of aperture functions and their corresponding PSFs. Figure (a) shows the formation of glare, revealing that the glare's area will expand with more dirt. Figure (b) shows the formation for different kinds of streaks. Figure (c) shows that the shimmer is mainly caused by polygon-shaped aperture or line-like scratches. Figure (d) represents the flares with the combination of these types of aperture dirt, suggesting that a flare image can be regarded as a sum of different components.

completely opaque. For a point High-Pressure Sodium (HPS) light, it can be viewed as a mixture of light with wavelength from 560mm to 630mm. Different typical types of aperture functions and their corresponding PSF are shown in Figure 2.

Based on the lens flare patterns, the aperture dirt function can be classified into three types: glare $\widetilde{A}_{gl}$, streak $\widetilde{A}_{st}$, and shimmer $\widetilde{A}_{sh}$. Thus, the Fourier transformation of aperture function $\widetilde{A}_{\lambda}(x, y)$ can be represented as the convolutions among these aperture dirt functions. Since all of these aperture dirt functions' Fourier transformations only have extremely high magnitude at the light source's position on the image plane, the PSF can be approximated as the sum of these functions' Fourier transformations:

$$PSF_{\lambda} \approx |\mathcal{F}\{A_a(x, y)\} + \mathcal{F}\{\widetilde{A}_{gl}(x, y)\} + \mathcal{F}\{\widetilde{A}_{st}(x, y)\} + \mathcal{F}\{\widetilde{A}_{sh}(x, y)\}|^2. \tag{13}$$

Each Fourier transformation in Equation (13) can be regarded as a typical type of flare components. Because the gamma correction after calculating the PSF can be approximated as a square root operation, the real-world lens flare image can be regarded as a sum of different types of flare components. As shown in Figure 2, the visual results of these compound flares show that our assumptions are quite reasonable.

We have introduced a simplified lens system without considering dispersion and transparent stains with depth. This simplified lens system can explain the physics of different common flare components. However, in a real-world situation, the diversity of aperture dirt and lens system makes it really difficult to design an aperture function that is as same as the type of target lens flare. Thus, we design a phenomenological method to synthesize lens flares rather than using a physics-based flare generation method. Our phenomenological method still produces reasonable flares.

2 Experimental details

Data augmentation To train our nighttime flare removal model, our paired flare-corrupted and flare-free images are generated on the fly. The flare-free images are sampled from the 24K Flickr images [12]. An inverse gamma correction with $\gamma \sim U(1.8, 2.2)$ is first applied to the flare image and flare-free image to recover the linear luminance. During the training stage, for each flare-free image, we randomly multiply the RGB values with $U(0.5, 1.2)$ and add a Gaussian noise with variance sampled from a scaled chi-square distribution $\sigma^2 \sim 0.01\chi^2$ to it. The result is then clipped to $[0, 1]$ as the corresponding flare-free image ground truth. For the image in our flare dataset, a randomly selected reflective flare is added to a random scattering flare to synthesize a flare image. A random rotation $U(0, 2\pi)$, random translation $U(-300, 300)$, random shear $U(-\pi/9, \pi/9)$, a random scale $U(0.8, 1.5)$, and a random flip are applied to this image. Then, a random weight in $U(0.8, 3)$ is multiplied with the brightness as color augmentation. Besides, a random blur with the blur size in $U(0.1, 3)$ and a global color offset in $U(-0.02, 0.02)$ can help solve the defocused situation and the situation in which the flare may illuminate the whole scene. Finally, the augmented flare and flare-free images will be randomly combined to create the flare-corrupted images.

Training details We follow Wu et al. [8] and use the same U-Net [4] as our nighttime flare removal baseline network model. During the training stage, our input flare-corrupted images are cropped to $512 \times 512 \times 3$ with a batch size of 2. We train our model for 60 epochs on 24K Flickr image dataset [12] with the Adam optimizer. The learning rate is 10^{-4}. We use the default setting for other parameters. During the training stage, the estimated flare image can be calculated with the subtraction of flare-corrupted image and flare-free image. The loss function is the combination of L_1 loss and perceptual loss with a pretrained VGG-19 [6]. After training, the saturated regions of a result are extracted and pasted back to the deflared image to recover the light source. After that, we achieve the final output. Since nighttime light sources are not always saturated in all three RGB channels, we set the threshold of luminance to 0.97 for saturation rather than 0.99 in Wu et al. [8]. Our strategy can effectively recover more tiny light sources.

To build a benchmark on our Flare7K dataset, we retrain the state-of-the-art image restoration networks including MPRNet [10], HINet [1], Uformer [7], and Restormer [9] using our Flare7K dataset. Among them, Uformer [7] and Restormer [9] adopt the Transformer-based structures. To ensure the fairness of the experiment, we use the same training setting and data augmentation method as stated above for training these methods. All these models are trained using the Nvidia Geforce RTX 3090 GPUs. Due to the limited GPU memory (24G memory) of Nvidia Geforce RTX 3090, we reduce the parameters of the MPRNet and Restormer which are known as the heavy networks. Specifically, we decrease the refinement block number of Restormer from 4 to 1 and set the dimension of the feature channel to 16 rather than the default dimension 48. The MPRNet's number of features is set to 24 rather than 40 to satisfy the memory limitation. We believe the default parameter settings of these two networks would performance better. The visual comparison results of different models on real-world nighttime flare images are shown in Figure 3.

3 More visual results

In this section, we show more visual comparisons on our synthetic test data and real test data. We set U-Net [4] as our baseline model and show the outputs between the baseline model trained on our data and prior works in Figure 4 and Figure 6. We also compare our baseline model trained on our data with the most related work by Wu et al. [8] on real-world images with different types of lenses in Figure 7.

To show the effectiveness of the flare removal algorithm trained on our flare dataset for downstream tasks, we compare the performance of the captured original images and processed images on the task of stereo matching. In this task, the apparent pixel difference for each corresponding pixel in stereo images is calculated to obtain the disparity map. In our experiment, we use a ZED camera to capture stereo images. Then, LEAStereo, a state-of-the-art stereo matching algorithm pre-trained on the Kitti dataset [3], is applied to these stereo images to estimate the disparity. As shown in Figure 5, it demonstrates that the flare removal algorithm trained on our dataset can boost the robustness of the stereo matching algorithm significantly.

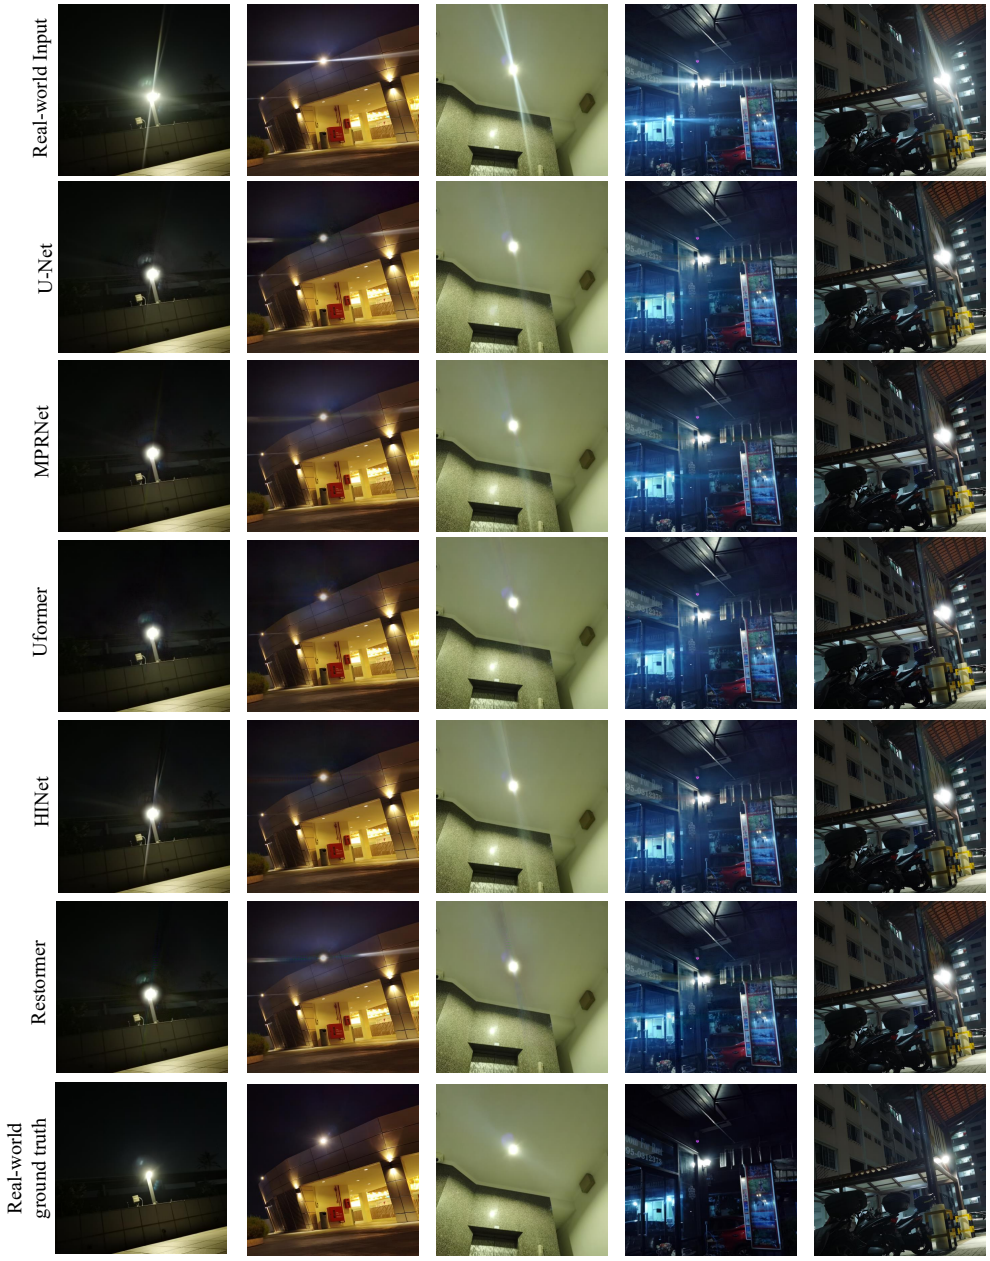

Figure 3: Visual comparison on our real-world test dataset for different image restoration networks trained on our Flare7K dataset. These networks include U-Net [4], MPRNet [10], Uformer [7], HINet [1], and Restormer [9].

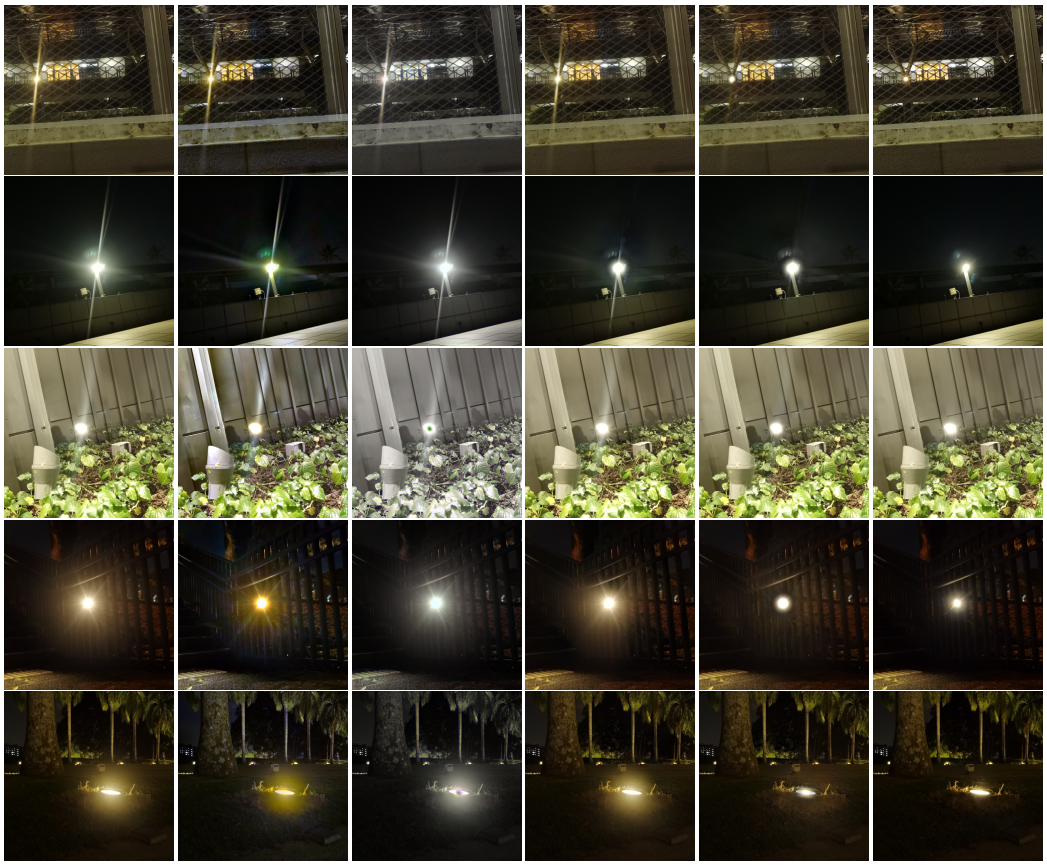

(a) Input (b) Zhang [11] (c) Sharma [5] (d) Wu [8] (e) Ours (f) Ground truth

Figure 4: Visual comparison on real-world nighttime flare images.

4 Extensions

In our dataset, for each lens flare image, we provide separated images of light source, glare with shimmer, streak, and reflective flare. These annotations can facilitate the design of improved flare removal methods and promote other related tasks. Besides, as the first nighttime flare dataset, the flare images can also be added to the training dataset of other nighttime vision algorithms to increase the robustness for flare-corrupted situations. With these annotations of our data, we provide more details about how to implement lens flare segmentation and light source extraction as follows. Please note that these annotations are not limited in these two applications.

Lens flare segmentation Flare segmentation is useful for flare removal. When the streak is too bright or overexposed, it is difficult to recover the details of these regions by just using an image decomposition network. A streak segmentation model trained on flare segmentation data can help locate these regions. Then, an image inpainting algorithm can be used to restore the missing information. Besides, the reflective flares for smartphone lenses always have the same patterns as the light sources' brightest regions. For matrix LED lights, the reflective flares will also be matrix-shaped patterns as shown in Figure 5 of the main paper. Thus, segmented reflective flares can be referred to restore the brightest regions in the saturated area, achieving the HDR snapshot reconstruction.

To demonstrate the effectiveness of flare segmentation, we train a network for lens flare segmentation and show the flare segmentation results in Figure 8. Specifically, we use the PSPNet[13] as the flare segmentation network and train it using the flare annotation of our dataset for 80k iterations with a batch size of 2 on an Nvidia GTX 1080 GPU. The data augmentation pipeline is the same as the pipeline mentioned in Section 2. We use the SGD optimizer with a learning rate 0.01, a weight decay

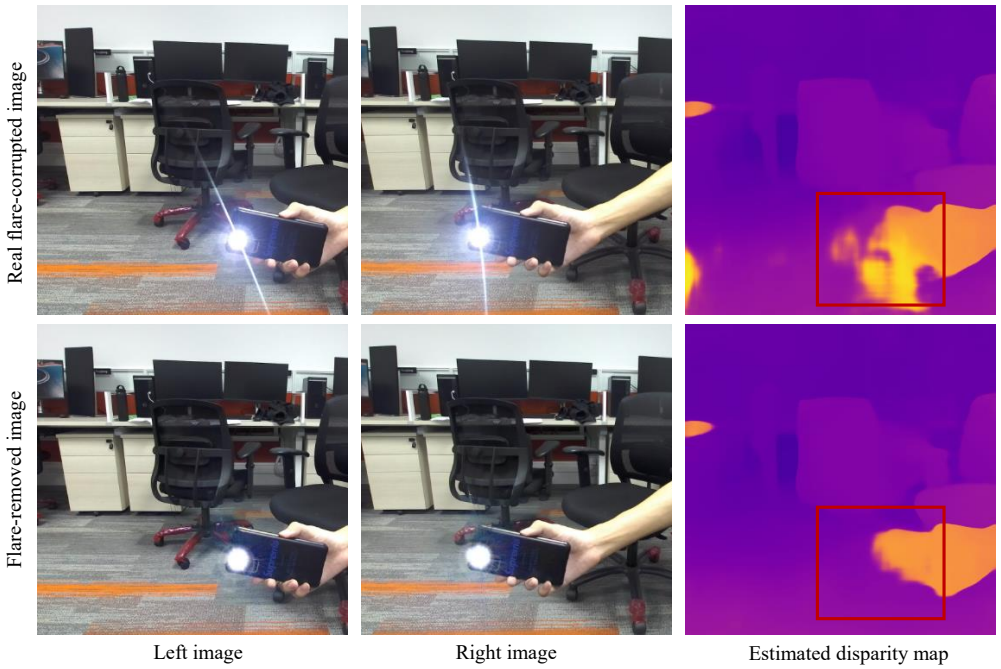

Figure 5: Visual comparison of estimated disparity map for flare-corrupted and flare-removed image pairs. The stereo images in the first row are captured by the ZED camera, and the second-row images are flare-removed images using deep model trained on our Flare7K dataset. Estimated disparity maps are calculated by using the LEAStereo [2] algorithm. As shown in the figure, the flare removal algorithm can be used to avoid mismatching caused by lens flare. The red boxes indicate the obvious differences.

0.0005, and a poly learning rate policy. The power rate of the policy is set to 0.9. Since areas of streak and light source are remarkably lower than the glare effect, we use a cross-entropy loss with the class weights 1.0, 1.0, 2.0, and 4.0 for background, glare with shimmer, streak, and light source.

Light source extraction Since flare removal algorithms may also remove the light source, as a part of the post-processing, one would need to extract the light source and then paste it back to obtain the final deflared image. As stated in Wu et al. [8], separating the light source from the flare-corrupted image is always challenging. Wu et al. use a feathered saturated mask to replace the light source. However, as shown in Figure 9(c), the light source processed by Wu et al. cannot simulate natural glow effect around the light source.

With the light source annotations provided by our dataset, paired scattering flare and light source can be used to train a light source extraction network. We use a U-Net as the light source extraction network and train it with L_1 loss on these 5000 scattering flare images for 50k iterations (20 epochs) on an Nvidia GTX 1080 GPU. Like our flare removal model, the batch size is set to 2. Moreover, we use the Adam optimizer with a learning rate 10^{-4}. For the flare removal model, the estimated flare can be calculated from the subtraction of the input and the de-flared images. Then, our light source extraction network can help separate the light source. Since the estimated flares do not have the same distribution as the training data, for data augmentation, we add low-frequency noise and use natural images with low luminance to the scattering flare as the training data. The steps above help improve the robustness of the light source extraction network. As shown in Figure 9(d), the glow effect of our method around the light source is more natural when compared with Wu et al. [8] as shown in Figure 9(c). It benefits from the high-quality annotations of the light source images in our scattering flares.

5 Social impacts

The proposed flare removal dataset is beneficial to increase the performance of a series of nighttime algorithms in flare-corrupted situations. Besides, it can enable more research on lens flare removal

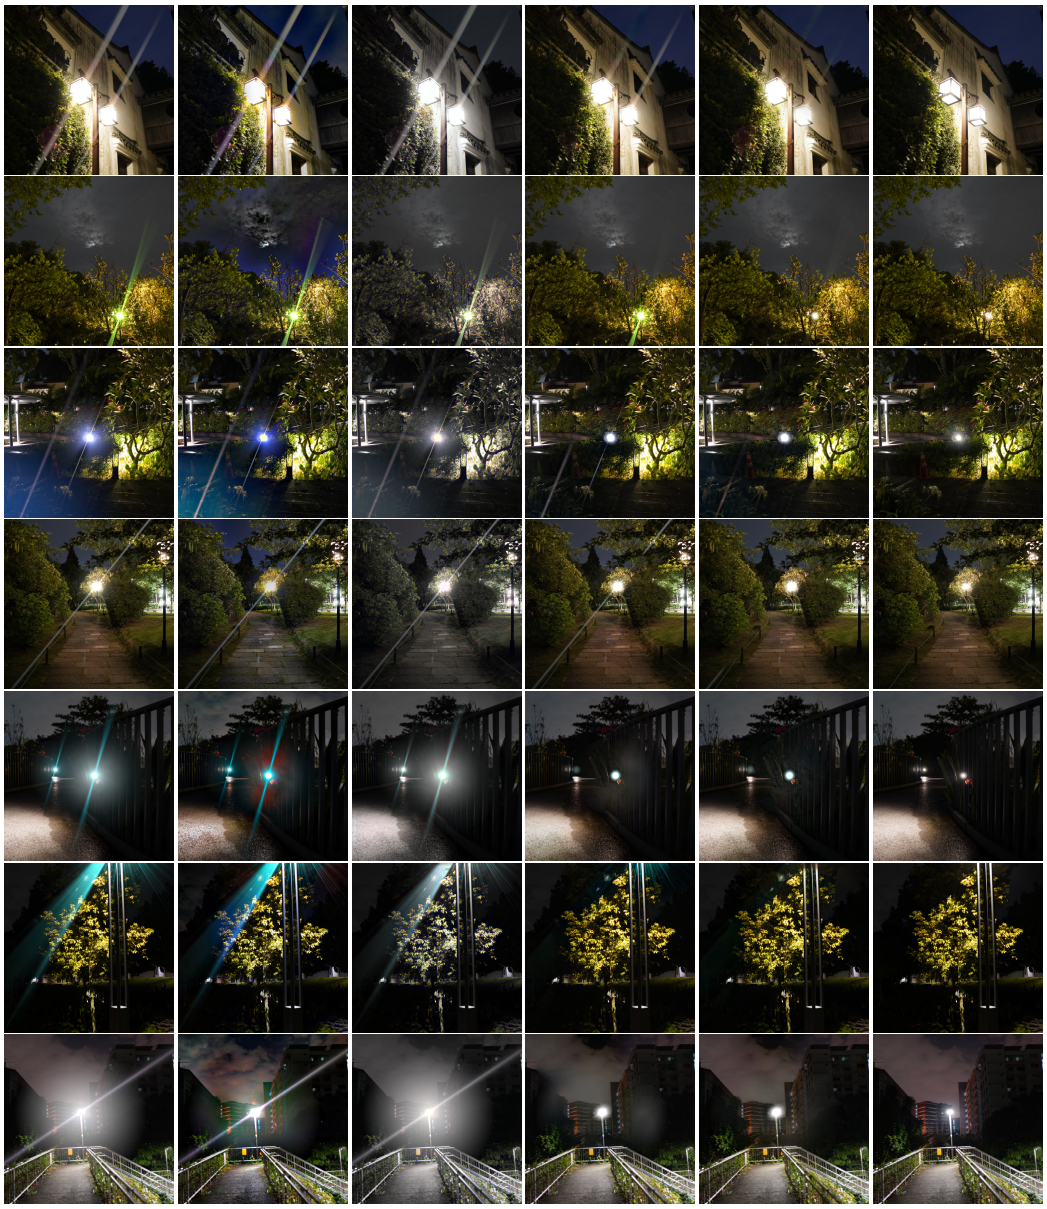

(a) Input (b) Zhang [11] (c) Sharma [5] (d) Wu [8] (e) Ours (f) Ground truth

Figure 6: Visual comparison on synthetic nighttime flare images.

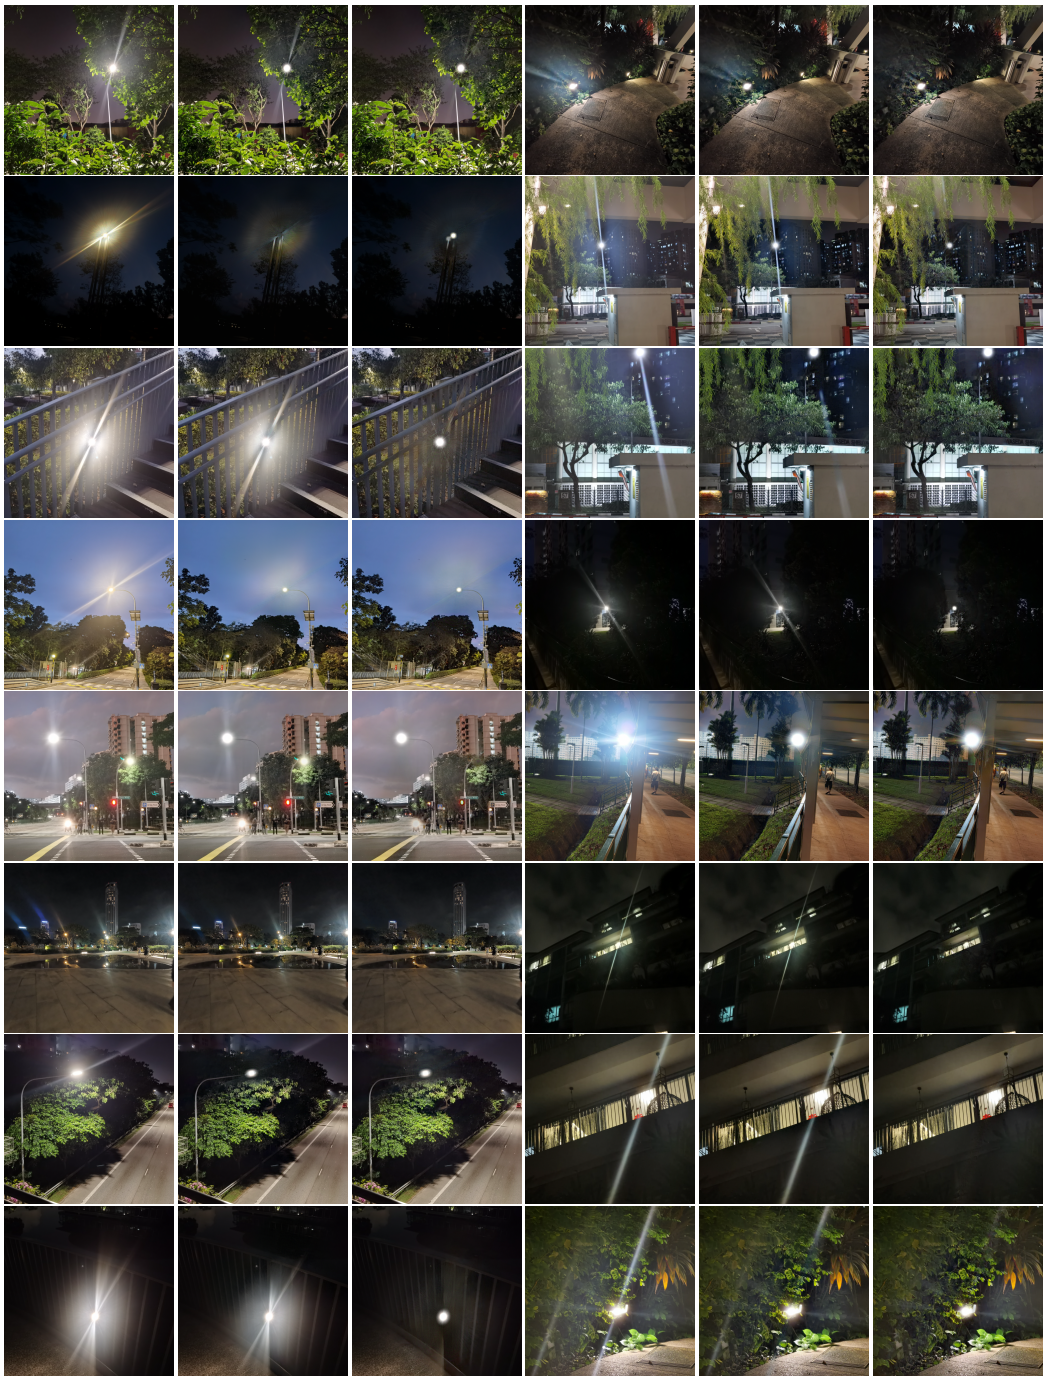

(a) Input (b) Wu [8] (c) Ours (d) Input (e) Wu [8] (f) Ours

Figure 7: Visual comparison on real-world flare images with different lenses and lights.

(a) Real input (b) Segmentation output (c) Real input (d) Segmentation output

Figure 8: Lens flare segmentation with our dataset. Since our dataset provides separated image for each component, it can be used for lens flare segmentation. This figure shows that PSPNet [13] trained on our dataset with annotations can accurately segment different components in real-world flare-corrupted images. In this figure, red, yellow, and blue represent streak, glare, and light source, respectively.

(a) Input (b) Estimated flare (c) Wu et al. [8] (d) Ours

Figure 9: Visual comparison between learning-based light source extraction method trained using our annotations and the feathered method proposed by Wu et al. [8] in real scenes. We use our flare removal model to separate the lens flare as shown in Figure (b). In Figure (c), the saturated region are processed by Wu et al. [8]' feathered method and pasted back to the de-flared image that is estimated by our flare removal model. In Figure (d), the light source extraction network stated in this section can provide more realistic glow effect around the light source than Wu et al. [8]' feathered method. (Zoom in for best view)

and nighttime visual enhancement. Our dataset also possesses many potential applications, such as nighttime automatic driving, license plate recognition, and nighttime object recognition. Since the dataset can simulate flare artifacts and degradation in the real world well, flares in our dataset can be added to the training set of other tasks as part of data augmentation. Such augmentation can improve the stability of the algorithms under different kinds of lens flare. Thus, we believe that our dataset will bring positive impacts on both academia and industry.

6 Dataset documentation

6.1 Motivation

- ***Q1****: For what purpose was the dataset created? Was there a specific task in mind? Was there a specific gap that needed to be filled?*

  A1: For commonly used cameras such as smartphone cameras and surveillance cameras, daily use may leave stains, scratches, and dust on lenses. While pointing to a strong light source, the scattering effect may lead to serious flare artifacts. This dataset is designed to help remove these flare artifacts at night. As stated in Section 3 of the main paper,

differences in light's spectrum and luminance bring an obvious domain gap between daytime and nighttime flares. As a result, existing synthetic flare dataset cannot perform well on the nighttime flare removal problem. To solve this gap, our Flare7K dataset is of great significance.

- ***Q2****: Who created the dataset (e.g., which team, research group) and on behalf of which entity (e.g., company, institution, organization)?*

  A2: This dataset is created by the team members of MMLab@NTU affiliated with S-Lab, Nanyang Technological University, Singapore.

6.2 Composition

- ***Q1****: What do the instances that comprise the dataset represent (e.g., documents, photos, people, countries)? Are there multiple types of instances (e.g., movies, users, and ratings; people and interactions between them; nodes and edges)? How many instances are there in total (of each type, if appropriate)? What data does each instance consist of?*

  A1: The training data of our Flare7k dataset contains 5,000 scattering flare images and 2,000 reflective flare images, consisting of 25 types of scattering flares and 10 types of reflective flares. Each reflective flare image can be added to a scattering image to compose a flare instance. For each scattering flare image instance, we provide separated streak and light source images. Thus, the reflective flare, streak, and light source for each flare instance can be used as annotations.

- ***Q2****: Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set? If the dataset is a sample, then what is the larger set? Is the sample representative of the larger set (e.g., geographic coverage)? If so, please describe how this representativeness was validated/verified. If it is not representative of the larger set, please describe why not (e.g., to cover a more diverse range of instances, because instances were withheld or unavailable).*

  A2: In this paper, we collect hundreds of flare-corrupted photos, and synthesize flare images based on these photos. At last our Flare7k dataset contains 5,000 scattering flare images and 2,000 reflective flare images, consisting of 25 types of scattering flares and 10 types of reflective flares. However, our dataset still cannot cover all possible instances in the wild as the diversity of real-world scenes. We conducte extensive experiments to demonstrate that our dataset can represent most real-world cases.

- ***Q3****: Are relationships between individual instances made explicit (e.g., users' movie ratings, social network links)? If so, please describe how these relationships are made explicit.*

  A3:N/A.

- ***Q4****: Are there recommended data splits (e.g., training, development/validation, testing)? If so, please provide a description of these splits, explaining the rationale behind them.*

  A4: Yes, we provide captured and synthetic test data. We also provide the script of data processing.

- ***Q5****: Are there any errors, sources of noise, or redundancies in the dataset? If so, please provide a description.*

  A5: We do not find any errors while synthesizing this dataset. If we find, we will provide an erratum.

- ***Q6****: Is the dataset self-contained, or does it link to or otherwise rely on external resources (e.g., websites, tweets, other datasets)? If it links to or relies on external resources, a) are there guarantees that they will exist, and remain constant, over time; b) are there official archival versions of the complete dataset (i.e., including the external resources as they existed at the time the dataset was created); c) are there any restrictions (e.g., licenses, fees) associated with any of the external resources that might apply to a dataset consumer? Please provide descriptions of all external resources and any restrictions associated with them, as well as links or other access points, as appropriate.*

  A6: Yes, this dataset is self-contained and licensed under CC BY-NC-SA 4.0.

- ***Q7****: Does the dataset contain data that might be considered confidential (e.g., data that is protected by legal privilege or by doctorpatient confidentiality, data that includes the content of individuals' non-public communications)? If so, please provide a description.*

A7: No.

- ***Q8****: Does the dataset contain data that, if viewed directly, might be offensive, insulting, threatening, or might otherwise cause anxiety? If so, please describe why.*

  A8: No.

6.3 Collection process

- ***Q1****: How was the data associated with each instance acquired? Was the data directly observable (e.g., raw text, movie ratings), reported by subjects (e.g., survey responses), or indirectly inferred/derived from other data (e.g., part-of-speech tags, model-based guesses for age or language)? If the data was reported by subjects or indirectly inferred/derived from other data, was the data validated/verified? If so, please describe how.*

  A1: Since our dataset consists of images, the data is directly observable.

- ***Q2****: What mechanisms or procedures were used to collect the data (e.g., hardware apparatuses or sensors, manual human curation, software programs, software APIs)? How were these mechanisms or procedures validated?*

  A2: We use Adobe After Effect to synthesize this dataset. To generate realistic flare images, we use Adobe After Affect's third-party plugin: Video Copilot's Optical Flares and Real Glow.

- ***Q3****: If the dataset is a sample from a larger set, what was the sampling strategy (e.g., deterministic, probabilistic with specific sampling probabilities)?*

  A3: This dataset is not a sample from a larger set.

- ***Q4****: Who was involved in the data collection process (e.g., students, crowdworkers, contractors) and how were they compensated (e.g., how much were crowdworkers paid)?*

  A4: This dataset is created by Yuekun Dai with the help of other co-authors. The data collection does not need compensation.

- ***Q5****: Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)? If not, please describe the timeframe in which the data associated with the instances was created.*

  A5: The flare-corrupted reference images were collected from November 2021 to April 2022. Our Flare7K dataset was created from April to June 2022. Our test set was collected in June 2022.

- ***Q6****: Were any ethical review processes conducted (e.g., by an institutional review board)? If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.*

  A6: No, there is no need for ethical review.

6.4 Preprocessing/cleaning/labeling

- ***Q1****: Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, tokenization, part-of-speech tagging, SIFT feature extraction, removal of instances, processing of missing values)? If so, please provide a description. If not, you may skip the remaining questions in this section.*

  A1: In our training set, our dataset is all synthetic. Thus, there is no preprocessing/cleaning/labeling. For our test dataset, we manually align the corresponding images as preprocessing.

- ***Q2****: Was the "raw" data saved in addition to the preprocessed/cleaned/labeled data (e.g., to support unanticipated future uses)? If so, please provide a link or other access point to the "raw" data.*

  A2: No, we will not provide raw misaligned test images.

- ***Q3****: Is the software that was used to preprocess/clean/label the data available? If so, please provide a link or other access point.*

  A3: We use Adobe After Effect to do the alignment. This software is publicly available.

6.5 Uses

- ***Q1****: Has the dataset been used for any tasks already? If so, please provide a description.*
  A1: No.
- ***Q2****: Is there a repository that links to any or all papers or systems that use the dataset? If so, please provide a link or other access point.*
  A2: No, it is a new dataset.
- ***Q3****: What (other) tasks could the dataset be used for?*
  A3: This dataset may be also used for flare classification and segmentation, low-light enhancement, and other downstream tasks which may suffer from serious lens flares.
- ***Q4****: Is there anything about the composition of the dataset or the way it was collected and preprocessed/cleaned/labeled that might impact future uses? For example, is there anything that a dataset consumer might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other risks or harms (e.g., legal risks, financial harms)? If so, please provide a description. Is there anything a dataset consumer could do to mitigate these risks or harms?*
  A4: N/A.
- ***Q5****: Are there tasks for which the dataset should not be used? If so, please provide a description.*
  A5: N/A.

6.6 Distribution

- ***Q1****: Will the dataset be distributed to third parties outside of the entity (e.g., company, institution, organization) on behalf of which the dataset was created? If so, please provide a description.*
  A1: No.
- ***Q2****: How will the dataset will be distributed (e.g., tarball on the website, API, GitHub)? Does the dataset have a digital object identifier (DOI)?*
  A2: This dataset will be distributed on website https://nukaliad.github.io/projects/Flare7K. It doesn't have a DOI.
- ***Q3****: Will the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)? If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.*
  A3: This dataset will be distributed under CC BY-NC-SA 4.0.
- ***Q4****: Have any third parties imposed IP-based or other restrictions on the data associated with the instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms, as well as any fees associated with these restrictions.*
  A4: No.
- ***Q5****: Do any export controls or other regulatory restrictions apply to the dataset or to individual instances? If so, please describe these restrictions, and provide a link or other access point to, or otherwise reproduce, any supporting documentation.*
  A5: No.

6.7 Maintenance

- ***Q1****: Who will be supporting/hosting/maintaining the dataset?*
  A1: This dataset will be maintained by this paper's authors, mainly by Yuekun Dai.
- ***Q2****: How can the owner/curator/manager of the dataset be contacted (e.g., email address)?*
  A2: The manager of the dataset can be contacted via YDAI005@e.ntu.edu.sg.
- ***Q3****: Is there an erratum? If so, please provide a link or other access point.*
  A3: No. If there exists an error, we will upload it on our website.

- ***Q4****: Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)? If so, please describe how often, by whom, and how updates will be communicated to dataset consumers (e.g., mailing list, GitHub)?*

  A4: This dataset may be updated to cover more types of flare patterns by this paper's first author. If so, updates will be announced on the website https://nukaliad.github.io/projects/Flare7K.

- ***Q5****: If the dataset relates to people, are there applicable limits on the retention of the data associated with the instances (e.g., were the individuals in question told that their data would be retained for a fixed period of time and then deleted)? If so, please describe these limits and explain how they will be enforced.*

  A5: This dataset is not related to people.

- ***Q6****: Will older versions of the dataset continue to be supported/hosted/maintained? If so, please describe how. If not, please describe how its obsolescence will be communicated to dataset consumers.*

  A6: When the dataset is updated to cover more types of flares. Older versions will be served as a subset of the whole dataset. Thus, the older dataset will still be supported.

- ***Q7****: If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so? If so, please provide a description. Will these contributions be validated/verified? If so, please describe how. If not, why not? Is there a process for communicating/distributing these contributions to dataset consumers? If so, please provide a description.*

  A7: As an important and new task, there are still many flare-corrupted cases that this dataset and existing methods cannot solve well. Thus, if anyone wants to contribute to the dataset, be free to use email or Github to contact us.

References

[1] Liangyu Chen, Xin Lu, Jie Zhang, Xiaojie Chu, and Chengpeng Chen. Hinet: Half instance normalization network for image restoration. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2021.

[2] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Hongdong Li, Tom Drummond, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *Advances in Neural Information Processing Systems*, 33:22158–22169, 2020.

[3] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.

[4] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2015.

[5] Aashish Sharma and Robby T. Tan. Nighttime visibility enhancement by increasing the dynamic range and suppression of light effects. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[6] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations*, 2014.

[7] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restorationn. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

[8] 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 *IEEE International Conference on Computer Vision*, 2021.

[9] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 2022.

[10] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2021.

[11] Jing Zhang, Yang Cao, Zhengjun Zha, and Dacheng Tao. Nighttime dehazing with a synthetic benchmark. In *ACM International Conference on Multimedia*, 2020.

[12] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[13] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2017.