

Dynamic Exposure Burst Image Restoration

Woohyeok Kim Jaesung Rim Daeyeon Kim Sunghyun Cho

POSTECH

Abstract

Burst image restoration aims to reconstruct a high-quality image from burst images, which are typically captured using manually designed exposure settings. Although these exposure settings significantly influence the final restoration performance, the problem of finding optimal exposure settings has been overlooked. In this paper, we present Dynamic Exposure Burst Image Restoration (DEBIR), a novel burst image restoration pipeline that enhances restoration quality by dynamically predicting exposure times tailored to the shooting environment. In our pipeline, Burst Auto-Exposure Network (BAENet) estimates the optimal exposure time for each burst image based on a preview image, as well as motion magnitude and gain. Subsequently, a burst image restoration network reconstructs a high-quality image from burst images captured using these optimal exposure times. For training, we introduce a differentiable burst simulator and a three-stage training strategy. Our experiments demonstrate that our pipeline achieves state-of-the-art restoration quality. Furthermore, we validate the effectiveness of our approach on a real-world camera system, demonstrating its practicality.

1. Introduction

Burst imaging is a powerful technique in digital photography that involves capturing a series of images in quick succession. This technique is particularly effective for high-quality image restoration in high-noise scenarios, such as low-light conditions, where high ISO settings are typically used, resulting in significant noise in the output images. Exploiting the random nature of noise, burst image restoration techniques capture multiple burst images and restore a single high-quality image from them. Due to its effectiveness, burst image restoration has been actively studied for decades and has become a key feature in modern cameras.

Recent advancements in burst image restoration primarily focus on improving alignment and fusion algorithms [4, 6, 7, 12, 13, 18, 22, 27, 45, 57]. Despite these advancements, one crucial aspect of burst imaging has been largely

overlooked: finding the optimal exposure settings (i.e., ISO and exposure time) for burst images to maximize restoration quality. Most existing methods rely on manually designed exposure settings, typically assuming uniform exposure settings across all frames. However, this approach results in burst images with the same noise levels and similar degrees of blur, which limits performance due to insufficient complementary information. Some approaches [15, 55] have proposed using non-uniform exposure settings across burst images, but they often depend on predefined exposure settings, such as fixed exposure brackets. These predefined settings may not always be optimal for varying shooting environments. For example, when capturing a static scene without camera or object motion, longer exposure times are generally preferred to suppress noise. In contrast, shorter exposure times are more suitable in dynamic scenes with significant motion, as they help mitigate motion blur.

In this paper, we propose a novel burst image restoration pipeline, *Dynamic Exposure Burst Image Restoration (DEBIR)*, which produces a single clean RAW image from burst RAW images in low-light conditions. DEBIR enables effective burst image restoration by adaptively predicting an optimal exposure time for each burst image based on the shooting environment. To this end, DEBIR consists of a novel *Burst Auto-Exposure Network (BAENet)* and a burst image restoration network. BAENet determines the optimal exposure times for burst images, which maximize the restoration network performance, based on a preview image, current exposure settings, and motion information. The imaging system then captures burst images using these predicted exposure times, and the restoration network processes them to restore a clean, blur-free, and noise-free image.

However, training BAENet presents a significant challenge, as it necessitates a sequence of ground-truth exposure times for a given scene, which are difficult to obtain. Identifying this ground-truth exposure-time sequence involves comparing the burst image restoration results of burst images with different exposure times for the same scene across all possible combinations. This task is practically unfeasible because collecting burst images with varying exposure times and their corresponding ground-truth clean images, as

well as searching through all possible combinations, is far from straightforward.

To address this challenge, we also propose a novel *differentiable burst simulator* along with an effective three-stage alternating training strategy for BAENet. Specifically, our differentiable burst simulator takes a set of exposure times as input and synthesizes burst images with realistic degradations, including noise and blur. The differentiable burst simulator is fully differentiable with respect to exposure times, enabling training of BAENet with respect to the restoration loss. During training, BAENet predicts the exposure times, and the differentiable burst simulator uses them to synthesize burst images. The burst image restoration network then processes these synthesized burst images to restore a clean image. We compute a loss between the restored image and the ground-truth clean image and update the weights of BAENet by backpropagating this loss. This training process eliminates the need for ground-truth exposure times, enabling effective training of BAENet.

We verify the effectiveness of our approach through detailed quantitative and qualitative analyses. Our contributions are summarized as follows:

- We propose *DEBIR*, a novel burst image restoration pipeline that predicts the optimal exposure times for burst images based on the shooting environment. To the best of our knowledge, our work is the first to predict optimal exposure time for each burst image by directly optimizing the exposure prediction network with the restoration loss.
- To effectively train *BAENet*, which predicts the exposure times for burst images, we propose a novel *differentiable burst simulator* along with an efficient training strategy.
- We show that *DEBIR* outperforms traditional exposure settings and removes the need for manual optimization, as validated through extensive experiments.

2. Related Work

Burst Imaging. Burst imaging has a wide range of applications, including denoising [12, 13], deblurring [1, 11, 27, 57], and super-resolution [4, 6, 7, 45]. Recently, numerous learning-based methods have been proposed. Bhat *et al.* [4] proposed DBSR that aligns burst images using optical flows and merges them using an attention-based fusion method. BIPNet [12] employs pseudo burst generation and implicit feature alignment using deformable convolution. Burstformer [13] adopts a transformer-based architecture for burst imaging. Mehta *et al.* [22] proposed a multi-resolution architecture and a transposed attention-based fusion method. However, these approaches only handle burst imaging with uniform exposure and focus on network design, rather than finding the optimal exposure settings.

Restoration using Non-uniform Exposures. Capturing multiple images with non-uniform exposures, e.g., using

exposure bracketing, has been widely used for HDR imaging [14, 19, 28, 32, 39, 43, 46, 48, 49, 54]. Beyond HDR imaging, several methods have explored non-uniform exposures for other tasks. One common strategy is dual-exposure imaging, which captures a pair of long- and short-exposure images and leverages their complementary information for denoising and deblurring [8, 17, 25, 37, 52, 53, 56]. Separately, Kim *et al.* [15] recently proposed a method to select the base frame among burst images captured with non-uniform exposures, while Zhang *et al.* [55] introduced an exposure bracketing approach for various restoration tasks. However, these non-uniform exposure approaches rely on fixed brackets, which may lead to suboptimal performance under varying scene dynamics and lighting conditions.

Auto-Exposure (AE). Most AE methods aim to adjust exposures to acquire images with proper brightness without considering subsequent image restoration processes [2, 3, 30, 36, 38, 40, 41, 44]. A few studies have explored learning-based AE for specific tasks, such as personalized imaging [51] and object detection [29], yet they typically assume single-image scenarios. Recently, Xu *et al.* [47] proposed a reinforcement learning-based HDR imaging method that sequentially predicts exposure settings for capturing input LDR images. However, this stepwise process lengthens the interval between images, making it time-consuming and unsuitable for burst imaging.

For restoration, approaches such as Digital-Gimbal [10], Active S-L [50], and Liba *et al.*'s method [18] have been proposed. Digital-Gimbal introduces learnable exposure parameters for burst images, which are optimized during training, but the parameters remain fixed after training regardless of the scene. Active S-L treats exposure prediction as a classification problem, selecting the exposure settings from a predefined set. However, this approach is not scalable, as the predefined set grows exponentially with the number of input images. Liba *et al.* employ motion metering to predict exposure times for burst imaging, but their approach assumes uniform exposure times. Unlike previous approaches, our method adaptively predicts exposure time for each burst image according to the shooting environment and seamlessly scales to an arbitrary number of images.

3. Image Degradation Model

The exposure e of an image I is determined by its exposure time t and gain g , as follows:

$$e = \frac{t \cdot g}{k}, \quad (1)$$

where k is a constant that depends on the specific camera. The exposure e determines the overall brightness of the image I , while the exposure time t and the gain g influence the level of blur and noise, respectively. Specifically, the effects



Figure 1. Overview of our pipeline. BAENet predicts the exposure times of each burst image from a preview image. Differentiable Burst Simulator generates burst images according to the exposure times. The restoration network then reconstructs a high-quality image from them. During inference, the simulator is removed, and the restoration network processes real burst images captured by our camera system.

of t and g on the image I in the camera RAW space can be described by an image degradation model defined as:

$$I = \text{quant} \circ \text{clip} \circ \text{cfa} \left(g \left(\int_0^t S_\tau d\tau + N_{\text{shot}} + N_{\text{read}} \right) \right), \quad (2)$$

where S_τ is the scene radiance at time τ in the camera RAW color space. N_{shot} and N_{read} are shot noise and read noise, respectively. $\text{quant}(\cdot)$, $\text{clip}(\cdot)$, and $\text{cfa}(\cdot)$ represent quantization, dynamic-range clipping, and color filter array (CFA) sampling, respectively. \circ indicates function composition. A long exposure time increases the signal-to-noise ratio (SNR), but may also cause blur in I as scene radiance changes over time. On the other hand, increasing the gain with short exposure time prevents blur, but amplifies noise.

Most existing burst image restoration approaches capture burst images with the same exposure time and gain. Consequently, the input burst images exhibit similar levels of noise and blur, making it challenging to leverage complementary information for effective burst image restoration. In contrast, DEBIR captures burst images with non-uniform exposure times and gains to effectively exploit complementary information from them. Furthermore, BAENet adaptively predicts exposure times according to the shooting environment, enabling more effective burst image restoration.

4. Dynamic Exposure Burst Image Restoration

Fig. 1 shows an overview of DEBIR. DEBIR assumes an imaging scenario in which live-streaming preview images are captured and displayed to the user prior to taking burst shots, a common feature in modern cameras. During the capture of preview images, the imaging system also performs AE to find a proper exposure e for the target scene.

Once the user presses the shutter button, BAENet predicts the optimal exposure times for n burst images, leveraging information derived from preview images. Specifically, for n burst images $\mathbf{I} = \{I_1, \dots, I_n\}$, BAENet predicts their optimal exposure times $\mathbf{t} = \{t_1, \dots, t_n\}$. The gains of the burst images $\mathbf{g} = \{g_1, \dots, g_n\}$ are then determined as $g_i = ke/t_i$ where e is the exposure value of the target scene estimated by AE during the preview stage. As a result, the burst images have the same brightness level but different levels of noise and blur to exploit the complementary information from them. Subsequently, our framework captures burst images in the camera RAW space using

the predicted exposure times and gains, and processes them with a burst image restoration network to produce a clean RAW image. During training, we simulate this process using a differentiable burst simulator in place of a real imaging system, enabling end-to-end optimization.

Our framework targets low-light conditions, where burst image restoration is most critical. Based on this, we assume that preview images have a constant exposure time, set to the maximum permissible exposure time for preview, while their gains vary depending on the light condition. In our experiments, we set the exposure time of preview images, t_p , to 1/120 seconds. In the following, we describe each component of DEBIR in detail.

4.1. BAENet

To predict optimal exposure times reflecting the shooting environment, BAENet takes a preview image I_p in the camera RAW space, its gain g_p , and motion magnitude m_p as inputs. We define the motion magnitude m_p as the average magnitude of the optical flow vectors estimated between I_p and its previous preview frame I'_p . For optical flow estimation, we convert the RAW I_p and I'_p to sRGB images by applying a simple ISP pipeline [5], and adopt an off-the-shelf optical flow estimator [42]. From I_p , g_p , and m_p , BAENet predicts n exposure times \mathbf{t} , from which their corresponding gains \mathbf{g} are computed using Eq. (1).

By leveraging the gain g_p and the motion magnitude m_p , BAENet adaptively predicts optimal exposure times based on the amount of noise, and on both camera motion and object motion in the target scene. Additionally, I_p enables BAENet to more accurately reflect the characteristics of the target scene, such as detailed distributions of noise and blur, which are not captured by g_p and m_p , as well as its content.

For the network architecture of BAENet, we adopt MobileNetV2 [35] due to its lightweight design and efficient computation. Specifically, for the input to BAENet, we construct a tensor by concatenating g_p and m_p to I_p along the channel dimension, and feed this tensor to BAENet. We modify the input channel size of MobileNetV2 accordingly. As g_p and m_p have large values on different scales, we normalize them so that they are in the range of $[0, 1]$ by simple shifting and scaling. We refer the reader to the supplementary material for more implementation details on BAENet.

Training BAENet to predict arbitrarily long exposure

times is challenging and unstable due to the unbounded search space. Thus, for stable training, we constrain the sum of the exposure times to be shorter than a predefined upper bound t_u , i.e., $\sum_i t_i \leq t_u$. To achieve this, we modify the last layer of MobileNetV2 to produce a vector of $n + 1$ dimensions and apply the bounded softmax function, ensuring that the sum of the output is 1, and that each element is positive. We then multiply the first n elements of the output vector by t_u to obtain the exposure times of the burst images. The $(n + 1)$ -th element of the output vector allows the sum of the n exposure times to be shorter than t_u , if necessary. Mathematically, t_i is obtained as:

$$t_i = t_u \cdot \text{softmax}_{\text{bounded}}(f_i, \epsilon), \quad (3)$$

where f_i is the i -th feature value in the last layer before the bounded softmax operation, and $\text{softmax}_{\text{bounded}}(f_i, \epsilon)$ is the bounded softmax function of f_i whose output is bounded into $[\epsilon, 1 - n\epsilon]$. We set $\epsilon = t_{\min}/t_u$ where t_{\min} is the minimum exposure time of our imaging system (1/240 sec.).

4.2. Burst Image Restoration Network

Once BAENet predicts optimal exposures for burst images, our imaging system captures burst images where each burst image I_i is a RAW image with the Bayer pattern. The burst image restoration network takes n RAW images as input and produces a single restored RAW image. For the architecture of the restoration network, we employ Burstormer [13], a state-of-the-art model for burst image restoration and enhancement tasks. Specifically, we adopt a variant of the Burstormer architecture modified for burst denoising¹. We adjust the channel sizes of the input layer and of the intermediate convolution layers, which are determined by the number of input images.

4.3. Differentiable Burst Simulator

To synthesize burst images with varying exposures and preview images while enabling end-to-end training of BAENet with the restoration loss, our training process employs a differentiable burst simulator, which is a module with no learnable parameters. Given exposure times \mathbf{t} , gains \mathbf{g} , and a scene radiance sequence \mathbf{S} , which will be defined later, the simulator generates burst images. To this end, it first computes the start time t_i^s and end time t_i^e of the exposure duration for each burst image I_i . Specifically, we calculate t_i^s and t_i^e as:

$$t_1^s = t^0, \quad t_i^s = t_{i-1}^e + \delta \text{ for } i > 1, \text{ and } t_i^e = t_i^s + t_i, \quad (4)$$

where t^0 is a constant to reserve a certain amount of time for preview images before the first burst image. δ represents the inter-frame gap between consecutive burst images, which we set $\delta = 7/1920$ sec. in our experiments.

Next, the simulator synthesizes each degraded RAW image I_i with realistic blur and noise corresponding to its exposure duration $[t_i^s, t_i^e]$ and gain g_i . Mathematically, the synthesis of each burst image I_i can be described as:

$$I_i = \text{syn}(\mathbf{S}, t_i^s, t_i^e, g_i), \quad (5)$$

where $\text{syn}(\cdot)$ represents a degraded image synthesis function, which is defined based on the image degradation model in Eq. (2). The scene radiance sequence \mathbf{S} models the varying scene radiance S_τ in Eq. (2) and is defined as a high-FPS video frame sequence in the camera RAW color space, i.e., $\mathbf{S} = \{S_0, S_1, \dots\}$. Each S_i in \mathbf{S} indicates scene radiance during a small exposure time e_S , which is set to 1/1920 sec. in our experiments.

To support back-propagation through the simulator, it must be differentiable with respect to the input exposure times \mathbf{t} . To achieve this, the degraded image synthesis function is designed to be differentiable with respect to t_i^s and t_i^e , supporting continuous values for t_i^s and t_i^e . Specifically, we define the degraded image synthesis function as:

$$\text{syn}(\mathbf{S}, t^s, t^e, g) = \text{clip} \circ \text{cfa}(S_{s,e} + gN), \quad (6)$$

where $S_{s,e}$ represents the integration of scene radiance during the exposure duration from t^s to t^e , and N models both photon shot noise and read noise. We define $S_{s,e}$ as:

$$S_{s,e} = \frac{1}{t^e - t^s} \left(\alpha_s S_{\lceil \bar{t}^s \rceil} + \sum_{\tau=\lceil \bar{t}^s \rceil}^{\lceil \bar{t}^e \rceil} S_\tau + \alpha_e S_{\lceil \bar{t}^e \rceil} \right), \quad (7)$$

where $\bar{t}^s = t^s/e_S$ and $\bar{t}^e = t^e/e_S$. α_s and α_e are blending weights for the first and last scene radiance frames defined as $\alpha_s = \lceil \bar{t}^s \rceil - \bar{t}^s$ and $\alpha_e = \bar{t}^e - \lceil \bar{t}^e \rceil$, respectively. These weights ensure smooth interpolation across frames, making $S_{s,e}$ differentiable with respect to the exposure times \mathbf{t} . For the noise N in Eq. (6), we model both photon shot noise and read noise using a heteroscedastic Gaussian distribution, following [55]:

$$N \sim \mathcal{N}(0, \lambda_{\text{read}} + \lambda_{\text{shot}} S_{s,e}), \quad (8)$$

where λ_{shot} and λ_{read} correspond to the shot noise and read noise parameters, respectively. However, directly implementing Eq. (8) is non-differentiable. To circumvent this, we adopt the reparameterization trick [16, 33]. Specifically, we reformulate Eq. (8) as $N = \sqrt{\lambda_{\text{read}} + \lambda_{\text{shot}} S_{s,e}} \cdot Z$, where $Z \sim \mathcal{N}(0, 1)$. This formulation renders the random variable N differentiable with respect to $S_{s,e}$, and ultimately with respect to \mathbf{t} . Details of the gain range and noise parameters are provided in the supplementary material.

Compared to Eq. (2), the degraded image synthesis function in Eq. (6) has two differences. Firstly, it does not include quantization to ensure non-zero gradients during

¹<https://github.com/akshaydudhanel6/Burstormer>

training. Secondly, it does not amplify scene radiance but amplifies only the noise N according to the gain g . This is because our framework assumes that input burst images have gains inversely proportional to their exposure times, which results in the same brightness level across the images while producing different noise levels.

Training of our framework requires not only burst images but also their corresponding ground-truth images and preview images. For the ground-truth image I_{gt} , we use the first sharp frame of I_1 , i.e., $I_{gt} = \text{cfa}(S_{\bar{t}^0})$. Regarding preview images I_p and I'_p , we synthesize them using the differentiable burst simulator as follows. First, we set the exposure durations $[t_p^s, t_p^e]$ and $[t_p'^s, t_p'^e]$ of I_p and I'_p as:

$$t_p^e = t^0 - \delta, \quad t_p^s = t_p^e - t_p, \quad (9)$$

$$t_p'^e = t_p^s - \delta_p, \quad t_p'^s = t_p'^e - t_p, \quad (10)$$

where δ_p represents the temporal gap between the two preview images, which we set 39/1920 seconds. We set t^0 to be a multiple of e_S that is greater than $2t_p + \delta_p + \delta$ to take into account the exposure times of the preview images and the temporal gaps, and to ensure that I_{gt} has sharp details and aligns well with I_1 . Then, using the differentiable burst simulator, we synthesize I_p and I'_p with their respective exposure durations $[t_p^s, t_p^e]$ and $[t_p'^s, t_p'^e]$ and the gain g_p , which is randomly sampled during training.

5. Training

While our differentiable burst simulator enables end-to-end training of both BAENet and the burst image restoration network, jointly training them from scratch is unstable and prone to local minima due to their mutual dependence. In early training, BAENet may initially predict random suboptimal exposure times, which can cause the restoration network to adapt prematurely to these exposure times. This early bias in the restoration network can then propagate back to BAENet, reinforcing exposure predictions that align with the network’s limited capabilities rather than the true optimal ones. To mitigate this issue, we adopt a three-stage strategy: (1) we first pre-train the restoration network to handle burst images with diverse exposure times, (2) we then train BAENet to predict the optimal exposure times using the pre-trained restoration network, and (3) we finally fine-tune the restoration network with BAENet to further enhance performance. In the following, we describe our training dataset and each stage in detail.

5.1. Datasets

To train our framework, we use a training dataset \mathcal{D} consisting of scene radiance sequences, i.e., $\mathcal{D} = \{\mathbf{S}_1, \mathbf{S}_2, \dots\}$. We synthesized our scene radiance dataset from a video dataset consisting of clean video clips. Specifically, for each video clip in the video dataset, we apply gamma expansion

to convert them into the linear sRGB space. Afterward, we convert the videos into the RAW color space by randomly sampling inverse color correction matrices (CCMs) and inverse white balance (WB) gains, and applying them to the videos. For the video dataset, we used video clips from the GoPro dataset [26], where each video clip is captured at 240 FPS and contains camera and object motions. By applying frame interpolation, we increase the frame rate by a factor of 8, resulting in video clips with a frame rate of 1920 FPS. Additionally, we also use the RealBlur dataset [34] to synthesize scene radiance sequences of static scenes. Additional dataset details are provided in the supplementary material. Through this process, we generated 5,219 scene radiance sequences for training and 532 for evaluation.

To prevent overfitting across training stages, we avoid reusing the same sequences for both the burst image restoration network and BAENet. Without this precaution, the burst image restoration network might become biased toward certain combinations of burst images and exposure times encountered during the first training stage. This bias could adversely affect BAENet in the second stage, causing it to favor those specific combinations and ultimately impairing its generalization capability. Thus, to mitigate this issue, we split the dataset \mathcal{D} into $\mathcal{D}_{\text{restore}}$ (4,092 sequences) and $\mathcal{D}_{\text{BAENet}}$ (1,127 sequences).

5.2. Pre-training Burst Image Restoration Network

Our training process starts with training the burst image restoration network. In this step, we train the restoration network using a loss $\mathcal{L}_{\text{restore}}$ defined as:

$$\mathcal{L}_{\text{restore}} = \|\text{res}_\phi(\mathbf{I}) - I_{\text{gt}}\|_1, \quad (11)$$

where $\text{res}_\phi(\cdot)$ is the restoration network parameterized by ϕ . \mathbf{I} is a burst image sequence synthesized using our differentiable burst simulator with randomly sampled exposure times and gains, and I_{gt} is the ground-truth image.

We synthesize the burst images using the following method. We first randomly sample a scene radiance sequence \mathbf{S} from $\mathcal{D}_{\text{restore}}$, and a gain g_p for the sampled sequence. Next, we randomly sample the exposure time t_i for the i -th burst image from $[t_{\min}, t_{\max}]$, where t_{\min} and t_{\max} are set to 1/240 sec. and 1/30 sec., respectively. We then set its gain g_i as $g_i = t_p \cdot g_p / t_i$, ensuring that all burst images in the sequence have the same brightness. Finally, using the sampled scene radiance sequence \mathbf{S} , and the sampled exposure times t_i and gains g_i , we synthesize burst images with the differentiable burst simulator.

5.3. Training BAENet

Once the burst image restoration network is trained, we train BAENet using the restoration network on $\mathcal{D}_{\text{BAENet}}$. We train BAENet in two steps, a warm-up step and a main training

step, to avoid local minima and to achieve high-quality exposure time predictions.

The warm-up step trains BAENet using a loss $\mathcal{L}_{\text{warm-up}}$:

$$\mathcal{L}_{\text{warm-up}} = \|\text{bae}_{\theta}(I_p, g_p, m_p) - \mathbf{t}_{\text{pseudo-gt}}\|_1, \quad (12)$$

where $\text{bae}_{\theta}(\cdot)$ represents BAENet parameterized by θ , and $\mathbf{t}_{\text{pseudo-gt}}$ is a vector with pseudo-ground-truth exposure times obtained by examining a predefined set of different combinations of exposure times. For the warm-up step, we generate a training dataset with pseudo-ground-truth exposure times as follows. First, we define a set $\mathcal{E} = \{\hat{\mathbf{t}}_1, \dots\}$ where $\hat{\mathbf{t}}_i$ is a predefined combination of exposure times. In our experiments, we define \mathcal{E} as $\mathcal{E} = \{(8, 8, 8, 8), (16, 16, 16, 16), (24, 24, 24, 24), (32, 32, 32, 32), (8, 16, 24, 32)\}/1920$ seconds. Next, we sample a scene radiance sequence \mathbf{S} from $\mathcal{D}_{\text{BAENet}}$ and synthesize the input to BAENet, i.e., a gain g_p , preview images I_p and I'_p , and motion magnitude m_p , as well as the ground-truth image I_{gt} . We then synthesize burst images for each of the exposure time combinations in \mathcal{E} , and perform burst image restoration to each burst image sequence, and compare their results to I_{gt} . Finally, we choose the exposure time combination of the best result as $\mathbf{t}_{\text{pseudo-gt}}$. Once the dataset is generated, we train BAENet using $\mathcal{L}_{\text{warm-up}}$.

After the warm-up step, the main training step trains BAENet using a loss $\mathcal{L}_{\text{DEBIR}}$:

$$\mathcal{L}_{\text{DEBIR}} = \|\text{res}_{\phi}(\text{sim}(\text{bae}_{\theta}(I_p, g_p, m_p))) - I_{\text{gt}}\|_1, \quad (13)$$

where $\text{sim}(\cdot)$ represents the differentiable burst simulator. In this step, we fix the restoration network, update only BAENet for stable training of BAENet.

5.4. Fine-tuning Burst Image Restoration Network

Once BAENet is trained, we fine-tune the burst image restoration network to further improve its performance for the predicted exposure times by BAENet. To this end, we minimize $\mathcal{L}_{\text{DEBIR}}$ in Eq. (13) with respect to the parameters of the restoration network for $\mathcal{D}_{\text{restore}}$, while fixing BAENet.

6. Experiments

Implementation We implemented DEBIR using PyTorch [31]. We pre-trained the burst image restoration network for 500 epochs, trained BAENet for 100 epochs, and fine-tuned the burst network for 50 epochs. The learning rates for each stage were set to $3\text{e-}4$, $1\text{e-}7$, and $1\text{e-}5$, respectively, and were decayed to $1\text{e-}8$, $1\text{e-}8$, and $1\text{e-}7$ using a cosine annealing scheduler [20]. We used the AdamW optimizer [21], and stage 2 included a 35-epoch warm-up training. All experiments, except for the ablation on the number of burst images, were conducted with an upper bound of exposure time $t_u = 128/1920$ sec. and a burst size of $n = 4$, using training images of resolution

| Methods | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow |
|---------------------------|-----------------|-----------------|--------------------|
| Digital-Gimbal [10] | 33.87 | 0.9309 | 0.187 |
| Active S-L [50] | 33.89 | 0.9379 | 0.176 |
| Average AE [2] | 34.69 | 0.9484 | 0.157 |
| Gradient AE [38] | 34.86 | 0.9494 | 0.156 |
| Exposure Bracket [15, 55] | 35.04 | 0.9481 | 0.164 |
| DEBIR (Ours) | 35.32 | 0.9519 | 0.154 |

Table 1. Quantitative comparisons on our test set. The best and second-best results are in bold and blue, respectively.

256×256 . Training was conducted on a PC equipped with four GeForce RTX 3090 GPUs, using a batch size of 4.

6.1. Comparison with Other Methods

We compare DEBIR with existing AE and predefined-exposure methods. Each method either predicts exposure times or uses predefined ones, and the differentiable burst simulator generates the corresponding burst images, which are then restored by the restoration network. We evaluate classical and learning-based AE methods. The classical baselines, Average AE [2] and Gradient AE [38], determine exposure based on global brightness and local gradients. Since they produce a single exposure time, we apply the exposure to all burst images. For learning-based methods, we include Digital-Gimbal [10] and modify it by introducing learnable exposure parameters for each burst image. We also modify Active S-L [50], originally a two-image gain classifier, to estimate exposure times instead of gains. As a predefined-exposure baseline, we adopt classical exposure bracketing [15, 55] with increasing exposure times of $\{8, 24, 40, 56\}/1920$ seconds.

All methods are integrated with Burstormer [13] and are trained on our training set. For Active S-L, we integrated it with Burstormer for two images, as the method is not scalable beyond two images. To ensure fairness, we applied the same pre-training and fine-tuning steps to the restoration network for all methods.

Tab. 1 presents quantitative results, where DEBIR outperforms all baselines. Average AE and Gradient AE predict a single exposure time that is applied to all burst images, leading to similar noise and blur levels in the burst images, which ultimately limits their restoration performance. Predefined exposure bracketing performs better since it uses non-uniform exposures, but its fixed schedule still cannot adapt to the scene, leading to suboptimal results. Among learning-based methods, Digital-Gimbal is limited because its exposure parameters are fixed after training, while Active S-L underperforms due to its two-image design that cannot scale to longer bursts, both leading to insufficient performance gains. In contrast, DEBIR predicts per-frame exposure times conditioned on the shooting environment, better balancing noise and blur and exploiting stronger complementary information. Fig. 2 shows that DEBIR generates the most visually pleasing results.

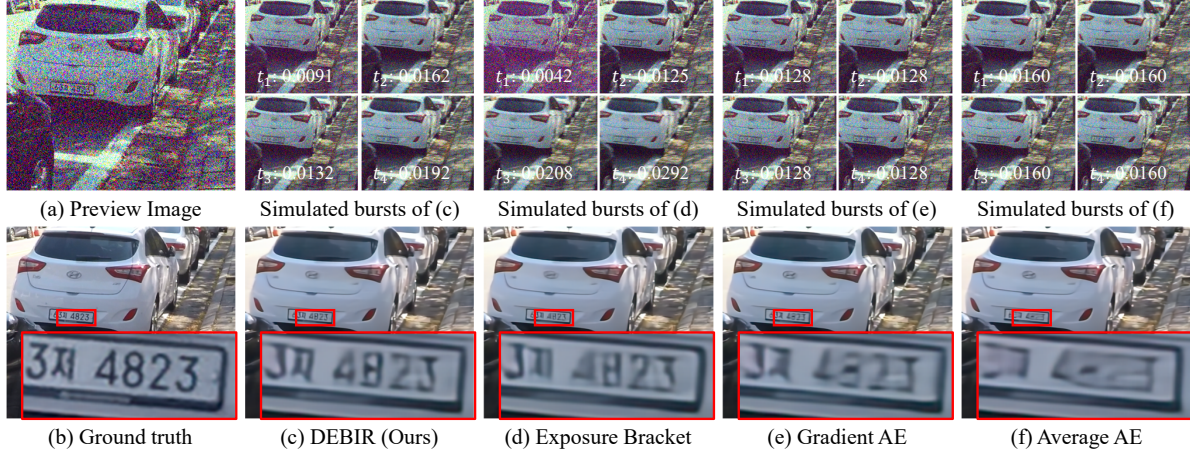


Figure 2. Qualitative results on our test set. See the supplementary material for more results.

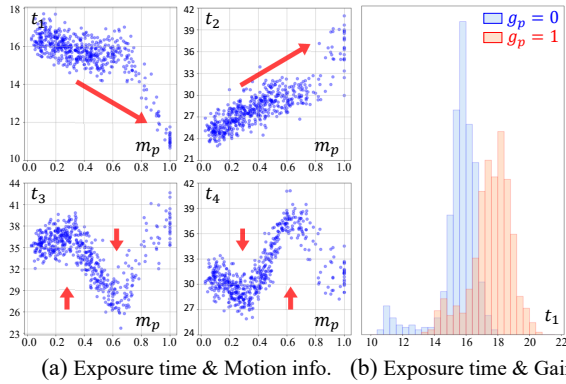


Figure 3. Analysis of predicted exposure times. (a) Scatter plots of exposure times of each frame t_i and motion magnitude m_p . (b) Exposure time histograms of the first frame (t_1) for the minimum and maximum preview image gains g_p . The unit of exposure times is $1/1920$ sec., and g_p is normalized to $[0, 1]$.

6.2. Analysis of BAENet

Predicted Exposure Time We analyze the predicted exposure time of each burst image t_i in relation to motion magnitude m_p and preview image gain g_p . Fig. 3-(a) and (b) show scatter plots of t_i versus m_p at $g_p = 0$ and histograms of t_1 at the minimum and the maximum g_p , respectively. The plots reveal two key observations. First, the exposure time of the first frame t_1 decreases as m_p increases to mitigate blur (Fig. 3-(a)). Additionally, increasing the gain g_p , which implies higher-noise conditions, shifts the predicted t_1 toward longer exposure times to suppress noise (Fig. 3-(b)). Second, BAENet optimizes the exposure times to maximize complementary information from adjacent burst images. For instance, in high-motion scenarios where m_p is large, BAENet decreases t_1 to minimize blur in the first frame. However, this also amplifies noise. To compensate for this, BAENet increases t_2 to capture complementary information, such as color details and noise-free information, which help mitigate noise in the first frame. A similar trend




Figure 4. Qualitative results using a real-world camera system.

| Methods | NIQE↓ [24] | BRISQUE↓ [23] | TOPIQ↑ [9] |
|------------------|------------|---------------|------------|
| Exposure Bracket | 6.57 | 50.66 | 0.339 |
| DEBIR (Ours) | 6.34 | 46.90 | 0.363 |

Table 2. Quantitative results using a real-world camera system.

is also found between the third and fourth frames, where t_4 adjusts in the opposite direction of t_3 (e.g., red arrows in Fig. 3-(a)). This result demonstrates the importance of complementary information in burst image restoration.

Real-world Capturing and Evaluation We capture real-world burst images using exposures predicted by BAENet, and compare the restored images to those from predefined-exposure settings [15, 55]. For a fair comparison, burst images from both methods must be captured simultaneously, as scene motion and content can change. To this end, we construct a handheld dual-camera system, depicted in Fig. 4-(a). The system consists of two cameras: one captures burst images using BAENet, while the other simultaneously captures burst images using exposure bracketing. We used the models trained on our dataset (Tab. 1). The inference time of BAENet, including optical flow estimation, is 0.023 seconds. Using the camera system, we captured and evaluated 142 burst images in low-light environments.

As ground-truth images are unavailable in practice, we evaluate the restored images using no-reference metrics. As shown in Fig. 4, the exposure bracket fails to restore accurate colors and produces artifacts. In contrast, our method successfully restores a high-quality image. Tab. 2 further validates the effectiveness of our method in real-world scenarios, highlighting its superior performance over prede-

| Inputs | w/o Prev. Image | w/o Gain | w/o Motion Info. | Full |
|-----------------|-----------------|--------------|------------------|--------------|
| PSNR \uparrow | 34.80 | 35.21 | 35.13 | 35.32 |

Table 3. Ablation study on the impact of each input in BAENet. The best and second-best results are in bold and blue, respectively.

| # of burst images | 2 | 4 | 6 | 8 |
|---------------------------|-------|-------|-------|--------------|
| Exposure Bracket [15, 55] | 33.59 | 35.04 | 35.69 | 35.89 |
| DEBIR (Ours) | 34.22 | 35.32 | 35.84 | 36.11 |

Table 4. Analysis of the impact of the number of burst images on performance in terms of PSNR. The best and second-best results are in bold and blue, respectively.

fined exposure settings. Additional results and details of our camera system are provided in the supplementary material.

6.3. Ablation Study

Uniform vs. Non-uniform Exposures In this study, we examine the impact of the non-uniform exposure setting in our method. To this end, we modify BAENet to predict a single exposure time and apply it uniformly across the burst. We then train the modified BAENet and the burst image restoration network using our training strategy. The variant achieves 35.01 dB PSNR, which is 0.31 dB lower than our final model. This result again shows the advantage of non-uniform exposure settings over uniform ones.

Inputs of BAENet BAENet takes a preview image, gain, and motion magnitude as inputs. We examine the impact of each input component. Tab. 3 shows that excluding any input significantly reduces performance, indicating that all inputs are crucial. Among them, the preview image has the most significant impact, as it contains all relevant information. However, providing motion magnitude and gain information helps quantify motion levels and noise, improving performance. Also, the table shows that the impact of motion information is more significant than that of gain.

Number of Burst Images While our experiments evaluate DEBIR with a fixed burst size of $n = 4$, the proposed method can be easily extended to other burst lengths. To verify this, we train DEBIR with $n = 2, 4, 6$, and 8, and evaluate its performance accordingly. For comparison, we also evaluate the performance of exposure bracketing methods [15, 55]. In these settings, the total exposure time is scaled proportionally to the burst size (i.e., 64, 128, 192, and 256 /1920 sec.), with the first exposure time fixed at 8/1920 sec. and subsequent exposure times increasing linearly (e.g., for $n = 4$, {8, 24, 40, 56}/1920 sec.). Tab. 4 shows that DEBIR outperforms the predefined exposure setting for all burst sizes. Performance improves with larger burst sizes, as more frames generally provide richer information.

Training Strategy We adopt a three-stage training strategy: pre-training the burst image restoration network (S1), training BAENet with $\mathcal{L}_{\text{warm-up}}$ (S2-1), training BAENet with $\mathcal{L}_{\text{DEBIR}}$ (S2-2), and fine-tuning the restoration network (S3).

| Training Strategy | PSNR \uparrow | SSIM \uparrow | LPIPS \downarrow |
|---------------------|-----------------|-----------------|--------------------|
| S1, S2-2 | 34.93 | 0.9482 | 0.162 |
| S1, S2-1, S2-2 | 35.01 | 0.9489 | 0.160 |
| S1, S2-2, S3 | 35.16 | 0.9502 | 0.157 |
| S1, S2-1, S2-2, S3 | 35.32 | 0.9519 | 0.154 |
| S1, S2-1, E2E | 34.99 | 0.9496 | 0.158 |
| S1, S2-1, S2-2, E2E | 35.11 | 0.9500 | 0.158 |
| E2E | 33.60 | 0.9279 | 0.196 |

Table 5. Ablation study on training strategy. S1: Pre-training the restoration network. S2-1: Training BAENet with $\mathcal{L}_{\text{warm-up}}$. S2-2: Training BAENet with $\mathcal{L}_{\text{DEBIR}}$. S3: Fine-tuning the restoration network. E2E: end-to-end training of BAENet and the restoration network. The best and second-best results are in bold and blue.

Tab. 5 shows the impact of each stage and comparisons with end-to-end training. The table shows that training BAENet with $\mathcal{L}_{\text{warm-up}}$ (S2-1) improves overall performance by providing pseudo-ground-truth exposure times in the warm-up stage. After training BAENet, fine-tuning the restoration network (S3) significantly improves the performance. This step further optimizes the network to adapt to optimal exposure settings estimated from BAENet. We also compare our approach with the end-to-end training strategy. Instead of fine-tuning the restoration network (S3), we apply end-to-end training. Additionally, we evaluate end-to-end training without the three-stage strategy. As shown in Tab. 5, end-to-end training leads to a significant performance drop, emphasizing the importance of our training strategy.

7. Conclusion

In this paper, we propose the Dynamic Exposure Burst Image Restoration (DEBIR) pipeline. Specifically, we introduce BAENet, which estimates exposure times to maximize complementary information across burst images. Then, a burst restoration network reconstructs high-quality images from burst images captured accordingly. For training them, we also present a novel Differentiable Burst Simulator and an efficient training strategy. We validate the effectiveness of DEBIR through extensive experiments. Finally, we integrate BAENet with a camera system and demonstrate the superior performance of DEBIR in real-world scenarios.

Limitations and Future Work BAENet may not always yield the optimal exposure settings for a given scene, as better choices could in principle be found through exhaustive search. This limitation is partly due to predicting exposure times for images yet to be captured with limited information and the risk of convergence to local minima. Addressing these issues presents an interesting direction for future work. Another promising future direction would be to extend our approach to jointly predict both exposure times and burst size, taking into account not only restoration quality but also practical constraints such as power consumption.

Acknowledgments This work was supported by grants from the Korea government (MSIT) through the National Research Foundation of Korea (NRF) (No. 2023R1A2C200494611) and the Institute of Information & Communications Technology Planning & Evaluation (IITP) (IITP-2026-RS-2024-00437866, No. RS-2019-II191906, Artificial Intelligence Graduate School Program (POSTECH)), as well as by the Basic Science Research Program through the NRF funded by the Ministry of Education (2022R1A6A1A03052954). It was also supported by Samsung Electronics Co., Ltd (IO251210-14286-01).

References

- [1] Miika Aittala and Frédo Durand. Burst image deblurring using permutation invariant convolutional neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [2] ARM. Mali-C71. <https://www.arm.com/products/silicon-ip-multimedia/image-signal-processor/mali-c71ae>, 2020. Camera product. 2, 6
- [3] Sebastiano Battiato, Arcangelo Ranieri Bruna, Giuseppe Messina, and Giovanni Puglisi. *Image Processing for Embedded Devices*. Bentham Science Publishers, 2010. 2
- [4] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Deep burst super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [5] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [6] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1, 2
- [7] Kelvin CK Chan, Shangchen Zhou, Xiangyu Xu, and Chen Change Loy. Basicvsr++: Improving video super-resolution with enhanced propagation and alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [8] Meng Chang, Huajun Feng, Zhihai Xu, and Qi Li. Low-light image restoration with short-and long-exposure raw pairs. *IEEE Transactions on Multimedia (TMM)*, 24:702–714, 2021. 2
- [9] Chaofeng Chen, Jiadi Mo, Jingwen Hou, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing (TIP)*, 33:2404–2418, 2024. 7
- [10] Omer Dahary, Matan Jacoby, and Alex M Bronstein. Digital gimbal: End-to-end deep image stabilization with learnable exposure times. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6
- [11] Mauricio Delbracio and Guillermo Sapiro. Burst deblurring: Removing camera shake through fourier burst accumulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [12] Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Burst image restoration and enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 1, 2
- [13] Akshay Dudhane, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Burstormer: Burst image restoration and enhancement transformer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2, 4, 6
- [14] Nima Khademi Kalantari, Ravi Ramamoorthi, et al. Deep high dynamic range imaging of dynamic scenes. *ACM Transactions on Graphics (TOG)*, 36(4):144–1, 2017. 2
- [15] Sanghyun Kim, Minjung Lee, Woohyeok Kim, Deunsol Jung, Jaesung Rim, Sunghyun Cho, and Minsu Cho. Burst image super-resolution with base frame selection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024. 1, 2, 6, 7, 8
- [16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014. 4
- [17] Wei-Sheng Lai, Yichang Shih, Lun-Cheng Chu, Xiaotong Wu, Sung-Fang Tsai, Michael Krainin, Deqing Sun, and Chia-Kai Liang. Face deblurring using dual camera fusion on mobile phones. *ACM Transactions on Graphics (TOG)*, 41(4):1–16, 2022. 2
- [18] Orly Liba, Kiran Murthy, Yun-Ta Tsai, Tim Brooks, Tianfan Xue, Nikhil Karnad, Qiuwei He, Jonathan T Barron, Dillon Sharlet, Ryan Geiss, et al. Handheld mobile photography in very low light. *ACM Transactions on Graphics (TOG)*, 38(6):164–1, 2019. 1, 2
- [19] Zhen Liu, Yinglong Wang, Bing Zeng, and Shuaicheng Liu. Ghost-free high dynamic range imaging with context-aware transformer. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [20] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017. 6
- [21] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 6
- [22] Nancy Mehta, Akshay Dudhane, Subrahmanyam Murala, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. Gated multi-resolution transfer network for burst restoration and enhancement. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 1, 2
- [23] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spa-

- tial domain. *IEEE Transactions on Image Processing (TIP)*, 2012. 7
- [24] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters (SPL)*, 2012. 7
- [25] Janne Mustaniemi, Juho Kannala, Jiri Matas, Simo Särkkä, and Janne Heikkilä. Lsd²: Joint denoising and deblurring of short and long exposure images with cnns. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2020. 2
- [26] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 5
- [27] Seungjun Nah, Sanghyun Son, and Kyoung Mu Lee. Recurrent neural networks with intra-frame iterations for video deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2
- [28] Yuzhen Niu, Jianbin Wu, Wenxi Liu, Wenzhong Guo, and Rynson WH Lau. Hdr-gan: Hdr image reconstruction from multi-exposed ldr images with large motions. *IEEE Transactions on Image Processing (TIP)*, 30:3885–3896, 2021. 2
- [29] Emmanuel Onzon, Fahim Mannan, and Felix Heide. Neural auto-exposure for high-dynamic range object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2
- [30] SangHyun Park, GyuWon Kim, and JaeWook Jeon. The method of auto exposure control for low-end digital camera. In *Proceedings of the International Conference on Advanced Communication Technology (ICACT)*, 2009. 2
- [31] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *Proceedings of the Neural Information Processing Systems Workshops (NeurIPSW)*, 2017. 6
- [32] K Ram Prabhakar, Rajat Arora, Adhitya Swaminathan, Kunal Pratap Singh, and R Venkatesh Babu. A fast, scalable, and reliable deghosting method for extreme exposure fusion. In *Proceedings of the IEEE International Conference on Computational Photography (ICCP)*, 2019. 2
- [33] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2014. 4
- [34] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 5
- [35] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [36] Simon Schulz, Marcus Grimm, and R Grigat. Using brightness histogram to perform optimum auto exposure. *WSEAS Transactions on Systems and Control*, 2(2):93, 2007. 2
- [37] Shayan Shekarforoush, Amanpreet Walia, Marcus A Brubaker, Konstantinos G Derpanis, and Alex Levinstein. Dual-camera joint deblurring-denoising. *arXiv preprint, arXiv:2309.08826*, 2023. 2
- [38] Inwook Shim, Tae-Hyun Oh, Joon-Young Lee, Jinwook Choi, Dong-Geol Choi, and In So Kweon. Gradient-based camera exposure control for outdoor mobile platforms. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)*, 29(6):1569–1583, 2018. 2, 6
- [39] Jou Won Song, Ye-In Park, Kyeongbo Kong, Jaeho Kwak, and Suk-Ju Kang. Selective transhdr: Transformer-based selective hdr imaging using ghost region mask. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 2
- [40] Yuanhang Su and C-C Jay Kuo. Fast and robust camera’s auto exposure control using convex or concave model. In *Proceedings of the IEEE International Conference on Consumer Electronics (ICCE)*, 2015. 2
- [41] Yuanhang Su, Joe Yuchieh Lin, and C-C Jay Kuo. A model-based approach to camera’s auto exposure control. *Journal of Visual Communication and Image Representation*, 36: 122–129, 2016. 2
- [42] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 3
- [43] Steven Tel, Zongwei Wu, Yulun Zhang, Barthélémy Heyrman, Cédric Demonceaux, Radu Timofte, and Dominique Ginjac. Alignment-free hdr deghosting with semantics consistent transformer. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2023. 2
- [44] Juan Torres and José Manuel Menéndez. Optimal camera exposure for video surveillance systems by predictive control of shutter speed, aperture, and gain. In *Real-Time Image and Video Processing*, 2015. 2
- [45] Renlong Wu, Zhilu Zhang, Shuohao Zhang, Hongzhi Zhang, and Wangmeng Zuo. Rbsr: Efficient and flexible recurrent network for burst super-resolution. In *Proceedings of the Chinese Conference on Pattern Recognition and Computer Vision (PRCV)*, 2023. 1, 2
- [46] Shangzhe Wu, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Deep high dynamic range imaging with large foreground motions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [47] Tianyi Xu, Fan Zhang, Boxin Shi, Tianfan Xue, and Yujin Wang. Adaptiveae: An adaptive exposure strategy for hdr capturing in dynamic scenes. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2025. 2
- [48] Qingsen Yan, Dong Gong, Qinfeng Shi, Anton van den Hengel, Chunhua Shen, Ian Reid, and Yanning Zhang. Attention-guided network for ghost-free high dynamic range imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [49] Qingsen Yan, Weiye Chen, Song Zhang, Yu Zhu, Jinqiu Sun, and Yanning Zhang. A unified hdr imaging method with pixel and patch level. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 2
- [50] Dan Yang, Samu Koskinen, and Joni-Kristian Kämäräinen. Active short-long exposure deblurring. In *Proceedings of*

the International Conference on Pattern Recognition (ICPR), 2022. [2](#), [6](#)

- [51] Huan Yang, Baoyuan Wang, Noranart Vesdapunt, Minyi Guo, and Sing Bing Kang. Personalized exposure control using adaptive metering and reinforcement learning. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 25(10):2953–2968, 2018. [2](#)
- [52] Lu Yuan, Jian Sun, Long Quan, and Heung-Yeung Shum. Image deblurring with blurred/noisy image pairs. *ACM Transactions on Graphics (TOG)*, 26(3):1–10, 2007. [2](#)
- [53] Zhilu Zhang, RongJian Xu, Ming Liu, Zifei Yan, and Wangmeng Zuo. Self-supervised image restoration with blurry and noisy pairs. In *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*, 2022. [2](#)
- [54] Zhilu Zhang, Haoyu Wang, Shuai Liu, Xiaotao Wang, Lei Lei, and Wangmeng Zuo. Self-supervised high dynamic range imaging with multi-exposure images in dynamic scenes. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2024. [2](#)
- [55] Zhilu Zhang, Shuohao Zhang, Renlong Wu, Zifei Yan, and Wangmeng Zuo. Exposure bracketing is all you need for a high-quality image. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025. [1](#), [2](#), [4](#), [6](#), [7](#), [8](#)
- [56] Yuzhi Zhao, Yongzhe Xu, Qiong Yan, Dingdong Yang, Xuehui Wang, and Lai-Man Po. D2hnet: Joint denoising and deblurring with hierarchical network for robust night image restoration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. [2](#)
- [57] Zhihang Zhong, Ye Gao, Yinqiang Zheng, and Bo Zheng. Efficient spatio-temporal recurrent neural network for video deblurring. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. [1](#), [2](#)