BlazeBVD: Make Scale-Time Equalization Great Again for Blind Video Deflickering

Xinmin Qiu¹, Congying Han¹, Zicheng Zhang¹, Bonan Li^{1 \boxtimes}, Tiande Guo¹, Pingyu Wang², and Xuecheng Nie³

¹ University of Chinese Academy of Sciences, Beijing, China ² Sichuan University, Sichuan, China ³ MT Lab, Meitu Inc., Beijing, China {qiuxinmin21,zhangzicheng19}@mails.ucas.ac.cn, {hancy,libonan,tdguo}@ucas.ac.cn, wangpingyu@scu.edu.cn, nxc@meitu.com



Fig. 1: Comparisons of the proposed *BlazeBVD*. We present flickering input, GT, Deflicker [21] and our BlazeBVD processed video frames, and illumination histograms along with KL divergence about GroundTruth. Our method recovers the illumination histograms well while avoiding the appearance of color artifacts and color distortions (such as the man's arm in the second column). Better see in color with $2 \times zoom$.

Abstract. Developing blind video deflickering (BVD) algorithms to enhance video temporal consistency, is gaining importance amid the flourish of image processing and video generation. However, the intricate nature of video data complicates the training of deep learning methods, leading to high resource consumption and instability, notably under severe lighting flicker. This underscores the critical need for a compact representation beyond pixel values to advance BVD research and applications. Inspired by the classic scale-time equalization (STE), our work introduces the histogram-assisted solution, called *BlazeBVD*, for high-fidelity and rapid BVD. Compared with STE, which directly corrects pixel values by temporally smoothing color histograms, BlazeBVD

 $^{^{\}boxtimes}$ Corresponding Author

leverages smoothed illumination histograms within STE filtering to ease the challenge of learning temporal data using neural networks. In technique, BlazeBVD begins by condensing pixel values into illumination histograms that precisely capture flickering and local exposure variations. These histograms are then smoothed to produce singular frames set, filtered illumination maps, and exposure maps. Resorting to these deflickering priors, BlazeBVD utilizes a 2D network to restore faithful and consistent texture impacted by lighting changes or localized exposure issues. BlazeBVD also incorporates a lightweight 3D network to amend slight temporal inconsistencies, avoiding the resource consumption issue. Comprehensive experiments on synthetic, real-world and generated videos, showcase the superior qualitative and quantitative results of BlazeBVD, achieving inference speeds up to $10 \times$ faster than state-of-the-arts.

Keywords: Video de-flicker · Histogram · Temporal consistency

1 Introduction

As social applications and the multimedia industry continue to expand, videos have become an essential medium for conveying information in daily life [30]. Achieving a high-quality video demands both image clarity and temporal consistency, nevertheless, video flickering frequently compromises the temporal consistency, often resulting from shooting environment and camera hardware limitations [9, 18]. This issue is further exacerbated when image processing techniques are applied to video frames. Moreover, flickering artifacts and color distortion are also prevalent in recent video generation technologies, including those based on Generative Adversarial Networks (GANs) [7,35,36] and Diffusion Models (DMs) [25, 42, 46, 52]. Since even certain types of flicker can significantly detract from the viewing experience, it is crucial to develop video deflickering techniques to eliminate flicker and preserve the integrity of video content given unknown degradation in various video processing scenarios.

Recovering visual content from unknown degradation is a key yet challenging topic [33, 43, 49]. Several seminal works [8, 20, 21] have addressed the blind video deflickering (BVD) task, in which methods refer to that (1) is agnostic to flickering patterns or levels (e.g., old movies, high-speed cameras, processing artifacts and color distortions), and (2) operates on a single flicker video without requiring additional guidance (e.g., flicker types, reference videos). In other words, BVD methods are blind to flicker types and guidance, making them widely applicable. Given these, the BVD task is challenging due to the lack of additional prior information guidance. Previous methods have explored traditional filtering [8,31], forced temporal consistency [20, 22], and atlas adjustment [21], etc.

Despite significant advancements by deep-learning methods in the BVD task. several critical issues hinder their broader application. First, these approaches demand substantial resources during both the training and inference stages, notably manifesting in slow processing speeds for individual videos. Second, their effectiveness is compromised by two main factors: (i) partial color artifacts and

 $\mathbf{2}$

distortion are present in the resulting video, causing differences in color fidelity and detail compared to the original video; (ii) Overexposure or underexposure due to lighting leads to significant loss of texture details and introduces new local flicker. These issues are largely attributed to the spatio-temporal complexity within video data, where the sheer volume of pixels presents a formidable obstacle for neural networks, including the advanced large vision models, to learn and maintain global visual consistency. Consequently, given that flicker fundamentally involves local or instantaneous shifts in illumination, effectively addressing the BVD task necessitates a representation more compact and adept at capturing illumination fluctuations than pixel values, whereas this aspect has not yet been explored by previous deep BVD methods.

In this paper, we draw inspiration from the classic flicker removal method scale-time equalization (STE) to facilitate deep BVD with histogram representation. An image histogram is defined as a distribution of pixel values, and it has been widely applied in image manipulation to adjust image brightness or contrast. Given an arbitrary video, STE temporally smooths histograms with Gaussian filtering and corrects pixel values in each frame with histogram equalization, improving the visual stability of the video. Although STE is only effective for some mild flickers, it validates that 1) histogram, which is much more compact than pixel values, can well delineate the illumination and flicker information, and 2) the video with a smoothed histogram series is visually appealing without obvious flicker, as shown in Fig. 1. Therefore, it is promising to advance the quality and speed of deep BVD with the hints from STE and histogram.

We introduce Blaze Blind Video Deflickering, dubbed as BlazeBVD, which is a histogram-assisted approach to achieve fast and faithful texture restoration given illumination fluctuation and over-/under-exposure. Compared to previous deep methods, BlazeBVD is the first to leverage histogram meticulously to ease the learning complexity of the BVD task. Specifically, BlazeBVD comprises three stages: At first, we introduce STE to rectify the histogram series of video frames under the illumination space, and extract *deflickering priors* including singular frames set, filtered illumination maps, and exposure maps. Second, due to the stable temporal performance of filtered illumination maps, they will be used as prompt conditions of a global flicker removal module (GFRM) containing 2D spatial networks to guide the color correction of video frames. On the other hand, a local flicker removal module (LFRM) is based on the optical flow warp to restore local over-/under-exposure regions labeled by exposure maps. Finally, we introduce a lightweight spatio-temporal network to process all frames, where an adaptive mask weighted warping loss is designed to improve video coherence. Our contributions can be summarized as follows:

(1) We present BlazeBVD, a histogram-assisted blind video deflickering method that simplifies the complexity and resource consumption of learning video data. At its core, BlazeBVD utilizes deflickering priors from STE, including filtered illumination maps for guiding the elimination of global flicker, singular frames set for identify the indexes of flicker frames, and exposure maps to identify regions affected locally by over-/under-exposure.

(2) Leveraging deflickering priors, BlazeBVD incorporates a Global Flicker Removal Module (GFRM) and a Local Flicker Removal Module (LFRM). These modules work together to efficiently correct global illumination and locally exposed textures in individual adjacent frames, significantly reducing processing time compared to handling the entire video. For enhanced coherence across frames, a lightweight spatio-temporal network is also integrated, boosting performance without significant time consumption.

(3) Through comprehensive experiments on synthetic, real-world and generated videos, we showcase the superior qualitative and quantitative results of BlazeBVD, achieving model inference speeds up to $10 \times$ faster than state-of-thearts. To the best of our knowledge, BlazeBVD also represents the first method to effectively tackle both illumination fluctuations and exposure challenges.

2 Related Work

Video deflickering addresses the degradation in temporal consistency [26,27, 29,38,48] resulting from abnormal camera, lighting, exposure or image processing algorithms. The relevant methods can be divided into three main categories. The first type of approach focuses on addressing specific flickering types, e.q., high-speed cameras [18], old film [9]. Delon [8] performs an axiomatic analysis of flicker, and proposed a fast method for global removal based on scale space theory to filter in color space directly. The second type of approach, blind temporal consistency [39, 53], aims to generate temporally stable output videos for arbitrary videos regardless of flickering or other artifacts. Bonneel et al. [4] calculate the gradient of the input frame as a guide to reduce the degree of inconsistency between frames. Lai et al. [20] input two consecutive frames as guidance. Lei etal. [22] directly learn the mapping function between input frames and processing frames to reach the purpose of inter-frame information consistency. Despite the generalization performance, such methods tend to be less efficient and hard to handle the complex flicker in practice. The third type of approach is blind video deflickering. Lei et al. [21] propose to solve universal blind video flicker using the neural atlas and video consistency. Meanwhile, some commercial software can remove flicker by integrating various known-flicker methods, such as DE-Flicker [34] and Flicker Free [2]. However, most of these methods require users to have flickering type knowledge background or slow repair speed. Our approach does not require prior cognition of the specific information flashing, and achieves fast deflickering speeds so that most users can efficiently process more videos.

Exposure correction. The purpose of exposure correction is to enhance an improperly exposed image to achieve a satisfactory visual effect. Information is lost due to over-/under-exposure input images. How to solve this problem intelligently and automatically, there have been many techniques in the image domain [12, 13, 15, 24, 41]. ERL [14] connects the under-/over-exposed optimization processes by correlating and constraining the relationships in the mini-batch correction process. DA [44] proposes to decouple contrast enhancement and detail recovery in each convolution process and design detail perception units to

be inserted into the existing CNN-based exposure correction network. However, there is no general approach regarding local exposure in video tasks.

Optical flow finds the corresponding relationship between the previous frame and the current frame by using the changes of pixels in the time domain and the correlation between adjacent frames, to calculate the motion information of objects between consecutive frames [10, 11, 16, 37]. In recent years, it has been used in video-related visual tasks [50]. The main target of flow-based methods is to produce accurate estimates of the optical flow through different network architectures, such as utilizing bidirectional optical flow estimation and refining inter-frame coherence.

3 Preliminaries

In this section, we introduce basic definitions and notations, and review relevant techniques involving image histograms for use in BlazeBVD.

Image histogram is a graphical representation illustrating the distribution of pixels across different intensity levels within an image [3]. For a single-channel image I with height H and width W, its histogram \mathcal{H} can be calculated as

$$\mathcal{H}(\lambda) = \frac{n_{\lambda}}{HW}, \quad if \ \lambda \in I \ otherwise \ 0, \tag{1}$$

where n_{λ} represents the count of pixels having the intensity λ .

Histogram matching or equalization [1, 5, 47, 51], is a technique to adjust the contrast of an image by transforming its histogram towards a predefined distribution. The cumulative histogram of an image histogram *hist* is defined by $Hist(\lambda) = \sum_{r=0}^{\lambda} \mathcal{H}(r)$. Give a target histogram \mathcal{H}' and its cumulative histogram $Hist'(\lambda)$, the histogram matching process can be described as

$$\mathcal{Q}(\lambda; \mathcal{H}, \mathcal{H}') = Hist'^{-1}(Hist(\lambda)), \ \forall \lambda \in I,$$
(2)

This procedure yields an image \tilde{X} , whose histogram closely aligns with *Hist'*. **Scale-Time Equalization** (STE) is a filtering approach [8] to adjust the intensity distribution across both spatial and temporal dimensions, mitigating flickers encountered in video sequences. The classic scale-space theory posits that convolving an image with a Gaussian kernel results in smoothing and the loss of image structure. Intuitively, while STE is based on scale-space theory, it applies Gaussian smoothing to sequences of histograms instead of the images. Given a flickering video sequence $\{X_t\}_{t=1}^T$, leveraging Eq. (1) can get the histogram sequence $\{\mathcal{H}_t\}_{t=1}^T$. For arbitrary time index t and intensity value λ , STE can be formulated as the convolution:

$$STE(\lambda, t) = G_s(\tau) * \mathcal{Q}(\lambda; \mathcal{H}_t, \mathcal{H}_\tau) \approx \sum_{\tau=t-l}^{t+l} G_s(\tau - t) \mathcal{Q}(\lambda; \mathcal{H}_t, \mathcal{H}_\tau), \qquad (3)$$

where *l* denotes the window radius, $G_s(t) = \frac{1}{\sqrt{4\pi s}} e^{-\frac{t^2}{4s}}$. After STE, we will get a sequence of histograms $\{\tilde{\mathcal{H}}_t\}_{t=1}^T$ and images $\{\tilde{X}_t\}_{t=1}^T$. Since histogram transformations do not create or cancel any image content, STE can achieve deflickering



Fig. 2: The framework of our approach *BlazeBVD*. We first extract flicker prior information about the input video and correct the brightness representation with STE-assisted histogram filtering in illuminance space (Stage1). Then, the prior is leveraged to remove temporal flicker and over-/under-exposure flicker from the global and local perspectives (GFRM and LFRM in Stage2). Finally, the temporal consistency of the processed video is improved (TCM in Stage3).

results without compromising the integrity of the original video. However, this also indicates that STE cannot compensate for the lost texture in flickering regions and produces vivid effects, which impedes its practical application.

4 Methodology

Let $\{X_t\}_{t=1}^T$ represent a video sequence affected by flickering defects and $\{G_t\}_{t=1}^T$ represent the groundtruth clean video. In this section, we propose *BlazeBVD* as a universal solution to eliminate these flicker and produce a temporally consistent video $\{O_t\}_{t=1}^T$. It is important to note that the types of flicker addressed in our work can vary widely, involving long-/short-term flicker in color and brightness.

Distinguished from existing prior BVD approaches, BlazeBVD employs a histogram-assisted approach, leveraging deflickering priors from Scale-Time Equalization (STE) [8] to streamline the complexity and resource demands of BVD task. Beyond ensuring stable temporal performance, BlazeBVD stands out as the first method capable of faithfully restoring the color and texture affected by illumination fluctuations and exposure challenges, including both over-/under-exposure.

4.1 Overview of BlazeBVD

As illustrated in Fig. 2, the BlazeBVD pipeline is structured into three main stages to address video deflickering challenges effectively: (1) Deflickering priors



Fig. 3: Pipeline of Stage1. The illumination maps are corrected in the temporal dimension by STE, which alleviates the difficulty of subsequent temporal data learning using networks. The deflickering priors can be extracted: \tilde{V}_t , M_t and $S_{flicker}$.

preparation. This initial stage leverages STE to correct the histogram within an illumination space representation. Key outputs from this stage include the filtered illumination map $\{\tilde{V}_t\}_{t=1}^T$, identifying the index set of singular frames $S_{flicker}$, and generating exposure maps $\{M_t\}_{t=1}^T$. These elements serve as crucial prior information, guiding the deflickering process in subsequent stages. (2) Global and local flicker removal. The second stage involves two flicker removal strategies. A global flicker removal module utilizes the filtered illumination map $\{\tilde{V}_t\}_{t=1}^T$ to address temporal flicker issues across the video. A local over-/underexposed region flicker removal module instructed by the singular frames set $S_{flicker}$ and exposure maps $\{M_t\}_{t=1}^T$, restores the texture in the over-/underexposure region to compensate for the lost high-frequency details. (3) Adaptive temporal consistency. The final stage introduces a spatio-temporal network under the guidance of an adaptive warping loss to improve video coherence and temporal consistency.

These stages delineated in the following Sec. 4.2, Sec. 4.3, and Sec. 4.4, together provide a comprehensive solution to enhance the visual quality and temporal consistency of video content.

4.2 Stage1: Deflickering priors preparation

Spatio-temporal challenge. An effective BVD model should strategically avoid the intricacies of spatio-temporal complexity across pixels, nevertheless, current deep learning approaches to BVD have not adequately addressed this challenge. For instance, the state-of-the-art BVD method, Deflicker [21], employs a neural atlas [19] as a unified representation of the entire video. While this can offer consistent guidance for the deflickering process, the effectiveness of neural atlases significantly depends on the video duration and the network capacity. Processing an 8-second video to create an atlas can take around 10 minutes, and longer videos exacerbate the demands on time, memory, and computational resources. Consequently, there is a further need for research into more compact representations that can facilitate deflickering without these limitations.

Illumination histogram. We introduce histogram representation [3], which utilizes statistical distribution to illustrate the global properties of image pixel values. As illustrated in Fig. 3, we first extract the illumination maps $\{V_t\}_{t=1}^T$ according to the conversion between HSV and RGB color space:

$$V_t = \max_{R,G,B \in X_t} \{R,G,B\}, \ t \in \{1,2,...,T\}.$$
(4)

Next, by applying Eq. (1) to each illumination map, we utilize illumination histograms $\{H_t\}_{t=1}^T$ to represent the illumination characteristics of video frames. These illumination histograms reflect important flickering information: As shown in Fig. 3, frames without distinct illumination changes have similar histograms, while the histogram of a flicker frame has a distribution distinguished from others. Furthermore, spikes in the head or tail area of the distribution indicate potential over-/under-exposure problems in the image. To this end, reasonable utilization of histogram information is significant for flicker elimination and alleviation of spatio-temporal complexity.

Deflickering priors. We introduce deflickering priors extracted from illumination histograms to facilitate the deflickering process. Initially, considering that the illumination histograms of natural videos exhibit continuous changes in the temporal dimension, we utilize Scale-Time Equalization (STE) as defined in Eq. (3) to smooth the illumination histograms, as illustrated in Fig. 3. Through histogram matching, we further obtain filtered illumination maps $\{\tilde{V}\}_{t=1}^T$ that exhibit no obvious flicker compared to the original ones, thereby enabling the subsequent independent processing of each frame by the neural network. Secondly, we identify the singular frames set $S_{flicker}$ containing the indexes of flicker frames, based on the divergence of histograms before and after STE:

$$S_{flicker} = \{t \in [1,T] | \mathrm{KL}(\tilde{H}_t \parallel H_t) > \bar{\epsilon}_t\}, \ \bar{\epsilon}_t = \frac{1}{2n+1} \sum_{\tau=t-n}^{t+n} \mathrm{KL}(\tilde{H}_\tau \parallel H_\tau), \ (5)$$

where n is the moving average radius. Finally, as shown in Fig. 2, we extract the exposure maps $\{M_t \in \mathbb{R}^{H \times W}\}_{t=1}^T$ of local over-/under-exposure according to the filtered illumination map:

$$M_t(i,j) = \begin{cases} 1 & \tilde{V}(i,j) < \epsilon_1 \text{ or } \tilde{V}(i,j) > \epsilon_2, \\ 0 & \text{otherwise.} \end{cases}$$
(6)

Here, ϵ_1 and ϵ_2 are darkness and exposure thresholds, respectively, and we consider illumination values belonging to $[\epsilon_1, \epsilon_2]$ as not posing an exposure risk.

In summary, we leverage illumination histograms to prepare deflickering priors including filtered illumination maps $\{\tilde{V}\}_{t=1}^T$, exposure maps $\{M_t\}_{t=1}^T$, and the singular frames set $S_{flicker}$, for alleviating the spatio-temporal complexity and enhancing deflickering in the subsequent learning stages.

4.3 Stage2: Global and local flicker removal

In this stage, simple 2D spatial neural networks are integrated with deflickering priors to eliminate both global and local flickering flaws. Here, "global" refers

to obvious temporal inconsistency due to illumination fluctuation, while "local" refers to lost texture details due to exposure and darkness. This approach allows for the direct generation of high-quality videos from flicker frames, regardless of the spatio-temporal complexity.

Global deflickering. BlazeBVD is equipped with a global flicker removal module (GFRM), which incorporates a 2D-Unet to correct each frame X_t with the guidance of the filtered illumination map \tilde{V}_t . In other words, the corrected image $\tilde{X}_t = Unet_{\theta}(X_t, \tilde{V}_t)$, where θ denotes the parameters of the 2D-Unet. The network is trained on the BVD dataset \mathcal{D} by minimizing the L2 loss function between the predicted images and the groundtruth images:

$$\mathcal{L}_t = \mathbb{E}_{(X_t, G_t) \in \mathcal{D}} \| Unet_\theta(X_t, \tilde{V}_t) - G_t \|.$$
(7)

Local deflickering. BlazeBVD is also equipped with a local flicker removal module (LFRM) to enhance local texture details. As depicted in Fig. 5, overexposure can transform intricate texture in local regions into uniform pure, while local under-exposure can result in black areas. This loss of detailed texture cannot be compensated for by previous BVD methods, such as Deflicker [21]. We leverage LFRM to transfer detailed information from neighboring frames to these ill regions. Given the corrected singular frame \tilde{X}_t in $S_{flicker}$ with local exposure, along with the previous frame \tilde{X}_{t-1} and the next frame \tilde{X}_{t+1} , we first compute the optical flow to extract part of the detailed texture:

$$o_{\tau \to t} = Flow(X_{\tau}, X_t), \ \tau \in \{t+1, t-1\} \text{ and } t \in S_{flicker}.$$
(8)

Here Flow is the optical flow estimation function, in our work we set it as the popular RAFT model [11, 37]. We migrate the possible detail textures of the preceding and following frames into the current frame:

$$\hat{X}_{t} = \mathcal{F}(M_{t} \odot \mathcal{W}(\tilde{X}_{t-1}, o_{t-1 \to t}) + (1 - M_{t}) \odot \tilde{X}_{t}, \tilde{X}_{t},
M_{t} \odot \mathcal{W}(\tilde{X}_{t-1}, o_{t+1 \to t}) + (1 - M_{t}) \odot \tilde{X}_{t}), t \in S_{flicker},$$
(9)

where \mathcal{W} is the warping operator and \mathcal{F} is the fusion network from [41].

4.4 Stage3: Adaptive temporal consistency

The former stages have well addressed the main illumination fluctuation and exposure problems in BVD task. Nevertheless, a video may also have content inconsistencies to largely degrade the visual effect, and the former stage may also cause this. Therefore, We introduce the Temporal Consistency Model (TCM) to refine video $\{\hat{X}_t\}_{t=1}^T$ as the final outputs $\{O_t\}_{t=1}^T$. TCM is developed based on the architecture by RTN [40], but with a temporal consistency loss of adaptive masks to refine local artifacts:

$$E'_{\text{pair}} (O_t, O_s) = \|W_t \odot (M_t + 1) \odot M_{t,s} \odot (O_t - W(O_s))\|_1,$$

$$\mathcal{L}_{warp} = \frac{1}{T - 1} \sum_{t=2}^T \left\{ E'_{\text{pair}} (O_t, O_1) + E'_{pair} (O_t, O_{t-1}) \right\},$$
 (10)

where W_t is the loss weight for the over-/under-exposure area and M_t is the mask obtained from the former stage. W is the warping operator and the frame after warping is computed by optical flow estimation [11,37]. $M_{t,s}$ represents the mask after optical flow estimation. This loss avoids the video blurring problem introduced in Deflicker [21]. The training loss of TCM consists of reconstruction loss \mathcal{L}_{rec} , perceptual loss \mathcal{L}_{per} [17], spatio-temporal adversarial loss \mathcal{L}_{adv} [6] and designed adaptive weighted warping loss \mathcal{L}_{warp} . Combining all the above losses, the overall objective to optimize TCM is:

$$\mathcal{L}_{TCM} = \alpha_1 * \mathcal{L}_{rec} + \alpha_2 * \mathcal{L}_{per} + \alpha_3 * \mathcal{L}_{adv} + \alpha_4 * \mathcal{L}_{warp}.$$
 (11)

In our work, we set $\alpha_1 = 1.0$, $\alpha_2 = 1.0$, $\alpha_3 = 0.01$ and $\alpha_4 = 0.1$.

5 Experiments

5.1 Datasets

We inherit the training dataset DAVIS-2017-Train, MS-COCO and the testing dataset *Blind Deflickering Dataset* from Deflicker [21], including Synthetic videos from DAVIS-2017-Val, Real-world videos and Generation videos.

Training dataset. MS-COCO [23] contains 118287 images in the train set and is used for the pre-training of Unet in GFRM. DAVIS-2017-Train [20, 28, 32] consists of 60 sequences with a total of 4219 labeled frames, with a video frame rate of 24fps and a resolution of 480p. We create the synthetic process that provides the underlying facts for quantitative analysis. The training dataset is degraded by performing flicker on DAVIS-2017-Train, as shown in Eq. (12):

$$X_t = G_t + F_t, \ t \in \{1, 2, ..., T\},\tag{12}$$

where $\{G_t\}_{t=1}^T$ are clean frames and $\{F_t\}_{t=1}^T$ are random synthesized flickering artifacts. The window size W in F_t has been set to denote the number of frames that share the same flickering artifacts, random sampling between [2, 12]. The value of W controls the degree of flicker in the short- and long-term of the video. **Testing dataset**. We conduct Synthetic videos, Real-world videos and Generation videos respectively, from the *Blind Deflickering Dataset* [21] and DAVIS-2017-Test processed by Eq. (12). Among them, Synthetic videos and DAVIS-2017-Test have ground truth and the others do not. We can only use no-referenced metric and qualitative observation. Generation videos refer to the processed videos generated by generative algorithms, which may lead to flicker artifacts due to the imperfections of the algorithms. Real-world videos including old movies, old anime, slow-motion, *etc.* are directly captured, unprocessed videos in the real world.

Comparison methods. We compare our BlazeBVD with other 4 methods, including STE [8], ConvLSTM [20], DVP [22] and Deflicker [21]. We set Deflicker as the baseline, since it is the first publicly presented method for BVD task and is currently the most advanced approach for general videos.



Fig. 4: Qualitative comparisons between previous methods and our *Blaze-BVD*. Our method removes flicker and restores details in over-exposed regions (*red boxes* in row 6 and row 8) while avoiding color artifacts (row 2 and row 4). Besides, *BlazeBVD* also ensures the fidelity of the video content and avoids color distortion (*cyan boxes* in row 5 and row 7). Zoom in for the best view and we recommend watching videos in the supplementary materials.

5.2 Comparisons with State-of-the-art Methods

Metrics. We quantitatively compare the differences between our method and others with reference and no-reference metrics, including PSNR, SSIM [45], and E_{warp} [22]. Among them, E_{warp} is a no-reference metric used in videos. For each frame O_t , we calculate the warping error with frame O_{t-1} and the first frame O_1 for considering both short-term and long-term consistency.

Quantitative results. We compare our BlazeBVD with the state-of-the-art method Deflicker [21], and other consistency methods. Table 1 and Table 2 provide the comparative experimental results of various methods on DAVIS-2017-Test and Synthetic videos, where PSNR and SSIM of the resulting video are higher than those of the other methods. This shows that BlazeBVD contributes to the fidelity maintenance of video content, as well as the correctness of the missing details complement. The warping error of our method is a little higher

Table 1: Quantitative comparisons of synthetic videos in DAVIS-2017-Test. Our method performs best in terms of PSNR and SSIM and E_{warp} of all the methods.

Method	RawVideo	$\operatorname{ConvLSTM}$	DVP	STE	Deflicker	Ours	GT
$PSNR\uparrow$	-	24.110	24.136	26.138	23.932	28.609	-
$\mathrm{SSIM}\uparrow$	-	0.9256	0.9263	0.9321	0.9243	0.9638	-
$E_{warp}\downarrow$	0.1429	0.1369	0.1415	0.1117	0.0840	0.0825	0.1097

Table 2: Quantitative comparisons in Synthetic videos from *Blind Deflick-ering Dataset*. W is the window size and represents the number of frames that share the same flicker artifact. L denotes the local seed window size.

	$E_{warp}\downarrow$				$\mathrm{PSNR}\uparrow$		$\mathrm{SSIM}\uparrow$	
videoiype	Raw	Deflicker	Ours	GT	Deflicker	Ours	Deflicker	Ours
Synthetic								
-W = 1	0.1933	0.0848	0.0895	0.1056	25.9312	29.6738	0.9403	0.9539
-W = 3	0.1522	0.0844	0.0890	0.1056	25.4996	28.8272	0.9380	0.9555
-W = 10	0.1307	0.0855	0.0904	0.1056	24.3842	26.9966	0.9344	0.9543
- $L = 3$	0.1113	0.0875	<u>0.0900</u>	0.1056	25.3874	27.0881	0.9271	0.9517
Average	0.1469	0.0856	<u>0.0897</u>	0.1056	25.3006	28.1464	0.9349	0.9538

than Deflicker, but it is closer to the ground truth. Through analysis, we find that warping error has inaccuracy in the estimation of optical flow, and this error is masked by image blurring in Deflicker. Therefore, we believe that E_{warp} , which has the same magnitude as baseline and GT, is the valid experimental comparison result. In Tab. 3, we compare BlazeBVD and Deflicker on Real-world videos and Generation videos. Since the videos in these two datasets with no reference are inherently discontinuous, where E_{warp} cannot directly represent video continuity, we present the comparative results of user studies. To evaluate the perceived preference between the baseline Deflicker and the proposed method, each user needs to select a video with better perceptual quality between these two choices. In total, there are 50 users and 50 pairs in 7 datasets of comparisons. Qualitative results. Both ConvLSTM and DVP are temporal consistency models that require a reference, importing the input video as the reference, and the processed video cannot remove significant flicker, as shown in Fig. 4. Additionally, the flicker caused by over-/under-exposure is prone to lose the local texture (column 1 in Fig. 4). The atlas-based representation in Deflicker cannot accurately locate these points with different local features, resulting in the loss of all details on the processed video. Instead, our method designs local positioning and adjacent frame details transfer in LFRM to preserve part of the texture on the road, as shown in the red box of Fig. 4. We also reveal the color artifact in column 4 of Fig. 4 that appears at baseline, due to the excessive weight of warping error in the training loss when refining temporal consistency. Therefore, our designed adaptive mask weighted training loss in TCM improves the temporal

Table 3: User studies of Real-world videos and Generation videos. Since the videos in these two datasets are inherently discontinuous, E_{warp} cannot properly reflect temporal consistency. Here we only present the comparative results of the user study.

Type	Expert	OldAnime	OldMovie	SlowMotion	TimeLapse	VideoDM	VideoLDM
$\operatorname{Raw}(E_{warp})$	0.0754	0.1045	0.1099	0.0724	0.0732	0.1248	0.1601
Deflicker Ours	38.67% 61.33%	19.25% 80.75%	22.83% 77.17%	37.67% 62.33%	31.43% 68.57%	30.50% 69.50%	26.34% 73.66%

consistency while avoiding image blurring. More visualization results are in the supplementary material.

5.3 Ablation Studies

STE filtering. We contrast the difference between introducing STE filtering in illumination space and directly in color space. As shown in the red box of Fig. 5, directly performing histogram correction in color channels separately introduces color artifacts and color distortion. Because the filtering of non-gray images in the three channels is relatively independent, the recombination after nonlinear transformation would produce such defects.

GFRM and LFRM. We compare the important roles of two modules, GFRM and LFRM. As shown in Tab. 5, removing GFRM and LFRM reduces PSNR and SSIM and increases E_{warp} of inference videos. It can be inferred that GFRM needs to initially remove global flicker between frames, otherwise subsequent modules cannot maintain the temporal consistency, while LFRM maintains the texture fidelity of the video content, which is reflected in the cyan box of Fig. 5. **Designed weighted loss**. As shown in Tab. 5, the lightweight spatio-temporal network TCM reduces E_{warp} , which refines the temporal consistency of processed videos. We also compare the effect of using unweighted training loss and the adaptive mask weighted loss on this temporal network. Figure 5 provides the resulting frame with unweighted training, which is generally blurred and suffers from color distortion, while the weighted training keeps the frame clear and faithful in most areas. Adaptive mask weighting improves the temporal consistency of the inference video, improving the inter-frame continuity more pertinently.

5.4 Discussion

Advantages and limitations. BlazeBVD is over $10 \times$ faster than the inference speed of baseline using the atlas representation, which is specifically recorded in Tab. 4. Our method also takes advantage of explicit prior extraction in the BVD task, which provides directions for improvement in the subsequent work. And the cause of local flicker has been noted and effectively dealt with from both global and local perspectives. However, our method still has drawbacks in LFRM. Due to the inaccurate optical flow motion estimation, the fusion network

Table 4: Inference time and model parameters comparison. Our approach is faster than the baseline in inference time with less memory, and comparable in MACs and Params.

Table 5: Ablation studies in DAVIS-2017-Test. w/o STE replaces STE with mean filtering in Stage1 and w/o weighted denotes the warping loss without mask weighting in TCM.

and Params			Method	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$E_{warp}\downarrow$
Method	Deflicker	Ours	w/o STE w/o GFRM	26.4412 23.8939	$0.9517 \\ 0.9244$	$0.1011 \\ 0.1225$
Time/80p Memory MACs	$\begin{array}{c} 614.19 \mathrm{s} \\ 5204 \mathrm{M} \\ 260.7 \mathrm{G} \end{array}$	58.37s 1389M 251.4G	w/o LFRM w/o TCM w/o weighted	27.1366 29.9130 28.1049	$0.9462 \\ 0.9583 \\ 0.9523$	$0.1122 \\ 0.1093 \\ 0.0849$
Params	12.48M	$17.77 \mathrm{M}$	BlazeBVD	28.6092	0.9638	0.0825



Fig. 5: Qualitative ablation studies of key designs in *BlazeBVD*. Specifically, local detail loss is represented in *cyan box*, color artifact is represented in the *red box*, temporal inconsistency is represented in the *black box*, and color distortion is represented in the *blue box*. Zoom in for the best view.

cannot convey the local texture of adjacent frames, and there are still slight edge artifacts and color distortion that cannot be eliminated. We hope to continue exploring the transfer and fusion of local region information between frames [40].

6 Conclusion

In this paper, we propose BlazeBVD, a universal approach for the blind video deflickering task. The core of our method is to prepare flicker priors within the STE filter in illumination space first, and then utilize these priors to correct the global flicker and local exposed texture. GFRM and LFRM leverage the priors to remove global and local flicker in several adjacent frames with less computation. Finally, TCM is used to improve video coherence and inter-frame consistency. Extensive experiments validate the effectiveness of the designed modules.

Acknowledgements

This paper is supported by the National Natural Science Foundation of China (Nos. U23B2012, 11991022), and the Fundamental Research Funds for the Central Universities (E3E41904X2).

References

- Afifi, M., Brubaker, M.A., Brown, M.S.: Histogan: Controlling colors of gangenerated and real images via color histograms. In: CVPR. pp. 7941–7950 (2021)
- 2. Anarchy, D.: Flicker free, https://digitalanarchy.com/Flicker/main.html
- Bassiou, N., Kotropoulos, C.: Color image histogram equalization by absolute discounting back-off. Computer Vision and Image Understanding 107, 108–122 (Jul 2007)
- Bonneel, N., Tompkin, J., Sunkavalli, K., Sun, D., Paris, S., Pfister, H.: Blind video temporal consistency. ACM Transactions on Graphics 34, 1–9 (Nov 2015)
- Bottenus, N., Byram, B.C., Hyun, D.: Histogram matching for visual ultrasound image comparison. IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control 68, 1487–1495 (May 2021)
- Chang, Y.L., Liu, Z.Y., Lee, K.Y., Hsu, W.: Free-form video inpainting with 3d gated convolution and temporal patchgan. In: ICCV. pp. 9066–9075 (2019)
- Chu, M., Xie, Y., Mayer, J., Leal-Taixé, L., Thuerey, N.: Learning temporal coherence via self-supervision for gan-based video generation. ACM Transactions on Graphics 39, 75–1 (2020)
- Delon, J.: Movie and video scale-time equalization application to flicker reduction. IEEE Transactions on Image Processing 15, 241–248 (Jan 2006)
- Delon, J., Desolneux, A.: Stabilization of flicker-like effects in image sequences through local contrast correction. SIAM Journal on Imaging Sciences 3, 703–734 (Jan 2010)
- Dosovitskiy, A., Fischer, P., Ilg, E., Hausser, P., Hazirbas, C., Golkov, V., Van Der Smagt, P., Cremers, D., Brox, T.: Flownet: Learning optical flow with convolutional networks. In: ICCV. pp. 2758–2766 (2015)
- 11. Eslami, N., Arefi, F., Mansourian, A.M., Kasaei, S.: Rethinking raft for efficient optical flow. arXiv preprint arXiv:2401.00833 (2024)
- Huang, J., Liu, Y., Fu, X., Zhou, M., Wang, Y., Zhao, F., Xiong, Z.: Exposure normalization and compensation for multiple-exposure correction. In: CVPR. pp. 6043–6052 (2022)
- Huang, J., Liu, Y., Zhao, F., Yan, K., Zhang, J., Huang, Y., Zhou, M., Xiong, Z.: Deep fourier-based exposure correction network with spatial-frequency interaction. In: ECCV. pp. 163–180 (2022)
- Huang, J., Zhao, F., Zhou, M., Xiao, J., Zheng, N., Zheng, K., Xiong, Z.: Learning sample relationship for exposure correction. In: CVPR. pp. 9904–9913 (2023)
- Huang, J., Zhou, M., Liu, Y., Yao, M., Zhao, F., Xiong, Z.: Exposure-consistency representation learning for exposure correction. In: ACM MM. pp. 6309–6317 (2022)
- Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., Brox, T.: Flownet 2.0: Evolution of optical flow estimation with deep networks. In: CVPR. pp. 2462–2470 (2017)

- 16 X. Qiu et al.
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV. pp. 694–711 (2016)
- Kanj, A., Talbot, H., Luparello, R.R.: Flicker removal and superpixel-based motion tracking for high speed videos. In: ICIP. pp. 245–249 (2017)
- Kasten, Y., Ofri, D., Wang, O., Dekel, T.: Layered neural atlases for consistent video editing. ACM Transactions on Graphics 40, 1–12 (2021)
- Lai, W.S., Huang, J.B., Wang, O., Shechtman, E., Yumer, E., Yang, M.H.: Learning blind video temporal consistency. In: ECCV. pp. 170–185 (2018)
- Lei, C., Ren, X., Zhang, Z., Chen, Q.: Blind video deflickering by neural filtering with a flawed atlas. In: CVPR. pp. 10439–10448 (2023)
- Lei, C., Xing, Y., Chen, Q.: Blind video temporal consistency via deep video prior. In: NeurIPS. pp. 1083–1093 (2020)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV. pp. 740–755 (2014)
- Ma, L., Ma, T., Liu, R., Fan, X., Luo, Z.: Toward fast, flexible, and robust low-light image enhancement. In: CVPR. pp. 5637–5646 (2022)
- Mei, K., Patel, V.: Vidm: Video implicit diffusion models. In: AAAI. pp. 9117–9125 (2023)
- Moniz, J.R.A., Kang, E., Póczos, B.: Luciddream: Controlled temporally-consistent deepdream on videos. arXiv preprint arXiv:1911.11960 (2019)
- Park, K., Woo, S., Kim, D., Cho, D., Kweon, I.S.: Preserving semantic and temporal consistency for unpaired video-to-video translation. In: ACM MM. pp. 1248–1257 (2019)
- Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: CVPR (Jun 2016)
- Pérez-Pellitero, E., Sajjadi, M.S., Hirsch, M., Schölkopf, B.: Perceptual video super resolution with enhanced temporal consistency. arXiv preprint arXiv:1807.07930 (2018)
- Pfeuffer, A., Dietmayer, K.: Separable convolutional lstms for faster video segmentation. In: ITSC. pp. 1072–1078 (2019)
- 31. Pitie, F., Kokaram, A., Dahyot, R.: Removing flicker from old movies. Master's thesis, University of Nice-Eophia Antipolis, France, Eeptember (2002)
- Pont-Tuset, J., Perazzi, F., Caelles, S., Arbeláez, P., Sorkine-Hornung, A., Van Gool, L.: The 2017 davis challenge on video object segmentation. arXiv:1704.00675 (2017)
- 33. Qiu, X., Han, C., Zhang, Z., Li, B., Guo, T., Nie, X.: Diffbfr: Bootstrapping diffusion model for blind face restoration. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 7785–7795 (2023)
- 34. RE:VISION: De:flicker, https://revisionfx.com/products/deflicker/
- 35. Saito, M., Saito, S., Koyama, M., Kobayashi, S.: Train sparsely, generate densely: Memory-efficient unsupervised training of high-resolution temporal gan. International Journal of Computer Vision p. 2586–2606 (Nov 2020)
- Skorokhodov, I., Tulyakov, S., Elhoseiny, M.: Stylegan-v: A continuous video generator with the price, image quality and perks of stylegan2. In: CVPR. pp. 3626–3636 (2022)
- Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: ECCV. pp. 402–419 (2020)
- Thasarathan, H., Nazeri, K., Ebrahimi, M.: Automatic temporally coherent video colorization. In: CRV. pp. 189–194 (2019)

- Thimonier, H., Despois, J., Kips, R., Perrot, M.: Learning long-term style preserving blind video temporal consistency. In: ICME. pp. 1–6 (2021)
- 40. Wan, Z., Zhang, B., Chen, D., Liao, J.: Bringing old films back to life. In: CVPR. pp. 17694–17703 (2022)
- Wang, H., Xu, K., Lau, R.W.: Local color distributions prior for image enhancement. In: ECCV. pp. 343–359 (2022)
- Wang, W., Liu, J., Lin, Z., Yan, J., Chen, S., Low, C., Hoang, T., Wu, J., Liew, J.H., Yan, H., et al.: Magicvideo-v2: Multi-stage high-aesthetic video generation. arXiv preprint arXiv:2401.04468 (2024)
- Wang, X., Li, Y., Zhang, H., Shan, Y.: Towards real-world blind face restoration with generative facial prior. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9168–9178 (2021)
- Wang, Y., Peng, L., Li, L., Cao, Y., Zha, Z.J.: Decoupling-and-aggregating for image exposure correction. In: CVPR. pp. 18115–18124 (2023)
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Transactions on Image Processing 13, 600–612 (2004)
- Wu, J.Z., Ge, Y., Wang, X., Lei, S.W., Gu, Y., Shi, Y., Hsu, W., Shan, Y., Qie, X., Shou, M.Z.: Tune-a-video: One-shot tuning of image diffusion models for textto-video generation. In: ICCV. vol. 39, pp. 7623–7633 (2023)
- 47. Xu, R., Tu, Z., Du, Y., Dong, X., Li, J., Meng, Z., Ma, J., Bovik, A., Yu, H.: Pik-fix: Restoring and colorizing old photos. In: WACV. pp. 1724–1734 (2023)
- Xu, X., Li, M., Sun, W., Yang, M.H.: Learning spatial and spatio-temporal pixel aggregations for image and video denoising. IEEE Transactions on Image Processing 29, 7153–7165 (2020)
- Yang, T., Ren, P., Xie, X., Zhang, L.: Gan prior embedded network for blind face restoration in the wild. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 672–681 (2021)
- Zhan, X., Pan, X., Liu, Z., Lin, D., Loy, C.C.: Self-supervised learning via conditional motion propagation. In: CVPR. pp. 1881–1889 (2019)
- Zhang, F., Shao, Y., Sun, Y., Zhu, K., Gao, C., Sang, N.: Unsupervised low-light image enhancement via histogram equalization prior. arXiv preprint arXiv:2112.01766 (2021)
- Zhang, Z., Li, B., Nie, X., Han, C., Guo, T., Liu, L.: Towards consistent video editing with text-to-image diffusion models. Advances in Neural Information Processing Systems 36 (2024)
- Zhou, Y., Xu, X., Shen, F., Gao, L., Lu, H., Shen, H.T.: Temporal denoising mask synthesis network for learning blind video temporal consistency. In: ACM MM. pp. 475–483 (2020)