

LAN-HDR: Luminance-based Alignment Network for High Dynamic Range Video Reconstruction

Haesoo Chung¹ and Nam Ik Cho^{1,2}

¹Department of Electrical and Computer Engineering, INMC, Seoul National University, Korea

²IPAI, Seoul National University, Korea
 {reneeish,nicho}@snu.ac.kr

Abstract

As demands for high-quality videos continue to rise, high-resolution and high-dynamic range (HDR) imaging techniques are drawing attention. To generate an HDR video from low dynamic range (LDR) images, one of the critical steps is the motion compensation between LDR frames, for which most existing works employed the optical flow algorithm. However, these methods suffer from flow estimation errors when saturation or complicated motions exist. In this paper, we propose an end-to-end HDR video composition framework, which aligns LDR frames in the feature space and then merges aligned features into an HDR frame, without relying on pixel-domain optical flow. Specifically, we propose a luminance-based alignment network for HDR (LAN-HDR) consisting of an alignment module and a hallucination module. The alignment module aligns a frame to the adjacent reference by evaluating luminance-based attention, excluding color information. The hallucination module generates sharp details, especially for washed-out areas due to saturation. The aligned and hallucinated features are then blended adaptively to complement each other. Finally, we merge the features to generate a final HDR frame. In training, we adopt a temporal loss, in addition to frame reconstruction losses, to enhance temporal consistency and thus reduce flickering. Extensive experiments demonstrate that our method performs better or comparable to state-of-the-art methods on several benchmarks. Codes are available at <https://github.com/haesoochung/LAN-HDR>.

1. Introduction

As diverse videos become easily accessible through video-on-demand services, demands for high-quality video content with high resolution and high dynamic range (HDR) are naturally increasing. HDR content can provide a rich

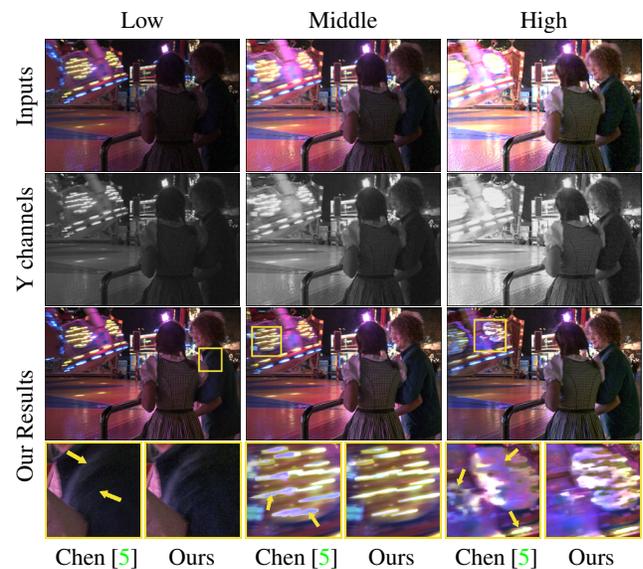


Figure 1: Visual comparisons of the proposed and Chen *et al.*'s method [5]. Our framework generates more detailed texture in regions with saturation and motions by utilizing luminance information.

viewing experience by displaying high contrast and a broad range of colors. While HDR displays are already ubiquitous, there is still a lack of HDR content available for delivery.

Whereas HDR imaging techniques for still images have been actively studied [8, 2, 33, 18, 45, 47, 35, 3, 38, 32, 10, 9, 1, 26, 49], those for videos have been relatively overlooked. Some early methods [41, 40, 23] proposed well-designed hardware systems for direct HDR video acquisition, but these systems mostly require a sophisticated design and high cost. Since Kang *et al.* [21] reconstructed an HDR video using a low dynamic range (LDR) image sequence with alternating exposures, this approach has been

commonly adopted. Specifically, an HDR frame is generated from a corresponding LDR frame and its neighboring frames, whose exposures alternate between two or three values, after motions between the frames are compensated. For instance, Mangiat *et al.* [28, 29] performed block-based motion estimation and then refined the results using color similarity and filtering, respectively. Kalantari *et al.* [20] synthesized multi-exposure images at each time step using patch-based optimization and merged them into an HDR frame. These approaches mostly require considerable time for the optimization process and suffer from artifacts resulting from inaccurate alignment.

With the development of deep learning, convolutional neural network (CNN)-based HDR video reconstruction methods have been proposed. Similar to the previous HDR imaging method [18], Kalantari *et al.* [19] aligned neighboring frames to the middle frame using optical flow and combined them with a simple merging network. More recently, Chen *et al.* [5] trained a two-step alignment network using optical flow and deformable convolution with large datasets. These methods significantly improved the quality of resulting HDR videos, but they still show some distortions and ghosting artifacts caused by optical flow estimation error, especially when motion exists in saturated regions or motion is fast. (See Fig. 1.)

In this paper, we propose a luminance-based alignment network (LAN) for HDR video composition. The LAN has a dual-path architecture, where an alignment module registers motions between frames using attention [42], and a hallucination module generates fine details, especially for saturated regions. In the alignment module, we rearrange a neighboring frame based on the attention score between itself and the reference frame. However, by applying naive attention, the system may find similar pixels solely by color. For content-based matching between two frames, we extract key and query features using only the luminance channel, which contains essential information on edge and structure. Here, downsampled frames are used as input to alleviate memory consumption. However, alignment cannot be perfect because downsampled inputs lack high-frequency information, and saturation hinders precise matching. Thus, the hallucination module fills in the missing details, especially in saturated areas using full-sized inputs, where we use gated convolution [50] with a mask that represents the image brightness. This mask enables the adaptive operation for very dark regions as well as highlighted areas. The features from each module are fused in an adaptive blending layer, which determines the contribution of each feature in terms of spatial dimension. Finally, we adopt a temporal loss for the resulting video to have consistent motions. Extensive experiments demonstrate that our method produces a clean HDR video from an LDR video with alternating exposures.

The main contributions of our paper are summarized as follows:

- We introduce an end-to-end HDR video reconstruction framework that performs a precise motion alignment using the proposed luminance-based alignment network (LAN). The LAN consists of two novel components: an alignment module and a hallucination module.
- The proposed alignment module performs content-based alignment by rearranging a neighboring frame feature (*value*) according to the attention score between content features from the neighboring frame (*key*) and the reference frame (*query*).
- The proposed hallucination module generates sharp details using adaptive convolution, especially for very bright or dark regions.
- We present a temporal loss to produce temporally coherent HDR videos without flickering.

2. Related Work

2.1. Multi-exposure HDR image reconstruction

HDR video reconstruction is similar to multi-exposure HDR imaging, a task of producing an HDR image from multiple differently exposed LDR inputs, in that they both aim to merge multi-exposure LDR frames with motions into a clean HDR frame. Most HDR imaging approaches focus on aligning input images to avoid ghosting artifacts. Some methods [22, 2, 15, 13, 51, 33, 24] assumed globally registered inputs and tried to detect and reject shifted pixels. Instead of simply discarding moving parts, several works performed more sophisticated motion compensation using optical flow [21, 52] and patch-based matching [39, 16]. However, most of these methods are slow due to the optimization process and often fail to handle large motions.

With the advancement of deep learning, HDR imaging networks have also been proposed. Kalantari *et al.* [18] and Wu *et al.* [45] proposed methods that align LDRs by flow-based motion compensation and homography transformation, respectively, and then merge them by simple CNNs. Pu *et al.* [37] presented a deformable convolution-based pyramidal alignment network, while Prabhakar *et al.* [36] trained a fusion network using both labeled and unlabeled data. Yan *et al.* [47, 48] aligned input images using spatial attention and non-local network, respectively. Note that these approaches differ significantly from ours in that they implicitly induce their attention maps to select valid areas through training. On the other hand, we compute the correlation between the key and query frames and then reconstruct aligned results based on the estimated scores. It is possible to apply the aforementioned methods to HDR video composition frame by frame, but temporal consistency is hardly guaranteed.

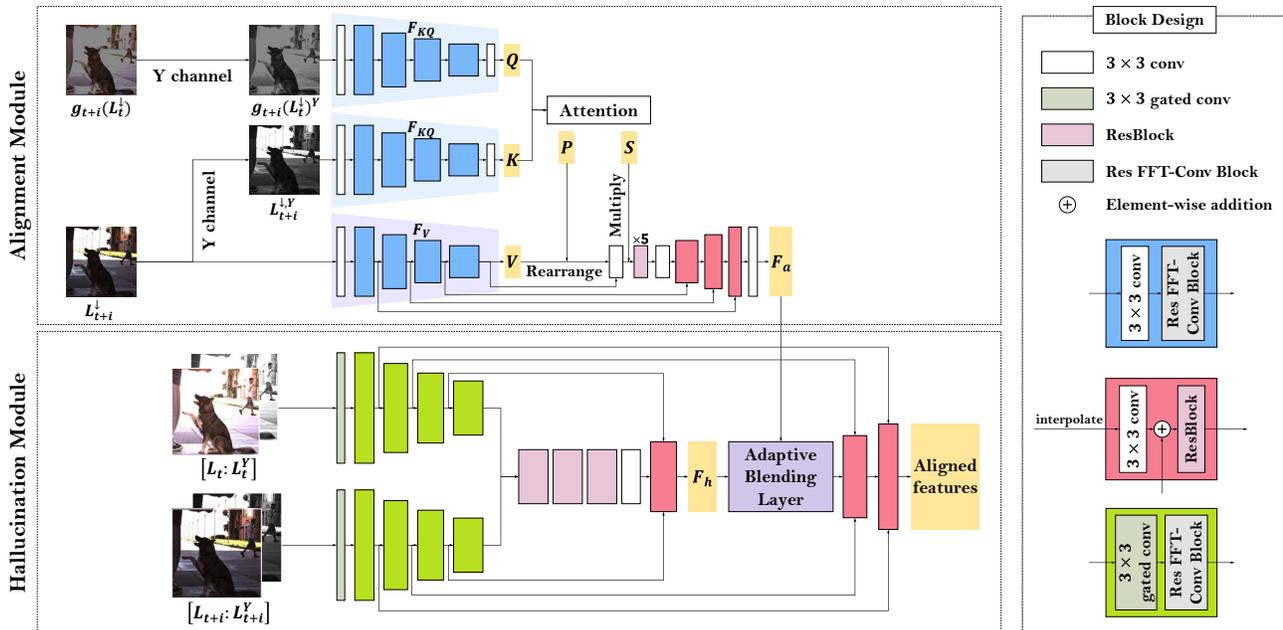


Figure 2: Overview of the proposed luminance-based alignment network (LAN).

2.2. HDR video reconstruction

Instead of directly acquiring HDR videos using specialized camera systems [41, 40, 23], which are very expensive approaches, many methods exploit LDR videos consisting of frames with spatially or temporally varying exposures. However, most approaches using dual- or tri-exposure images [7, 17] experience spatial artifacts. Meanwhile, similar to the multi-exposure HDR image fusion, most of the methods using a sequence with alternating exposures align multiple frames and combine them to produce an HDR frame. Kang *et al.* [21] proposed the first algorithm in this category, which synthesizes aligned multi-exposure images at each time step using global and local registration and merges them into an HDR frame. Kalantari *et al.* [20] also reconstructed missing images with different exposures based on patch-based optimization [39]. Mangiat *et al.* [28] adopted block-based motion estimation and additionally refined motion vector, which was improved by Mangiat *et al.* [29] with the addition of HDR filtering for block artifact removal. Meanwhile, Gryaditskaya *et al.* [14] proposed an adaptive metering algorithm to minimize motion artifacts. Li *et al.* [25] avoided correspondence estimation by formulating HDR video reconstruction as a maximum a posteriori (MAP) estimation problem, where foreground and background are synthesized separately and then combined together. These methods are mostly time-consuming and often display visible artifacts.

The recent work of Kalantari *et al.* [19] first presented a

CNN-based framework that consists of a flow network and a merging network, similar to the HDR imaging method in [18]. However, the flow estimation error often leads to undesirable artifacts in the resulting images. Chen *et al.* [5] trained a coarse-to-fine network, including a pyramid, cascading, and deformable alignment module [44] with newly collected datasets. This work requires global alignment as preprocessing and still uses optical flow for coarse alignment. Despite the introduction of various alignment strategies, these approaches often suffer from artifacts from misalignment in saturated areas and temporal inconsistency due to frame-by-frame reconstruction. In contrast, we perform luminance-based alignment without pre-alignment and improve temporal coherency by adding a temporal loss.

3. Proposed Method

Given an LDR video with alternating exposures, our goal is to generate a high-quality HDR video. In order to produce an HDR frame H_t , we use $2N + 1$ consecutive LDR frames $\{L_{t-N}, \dots, L_t, \dots, L_{t+N}\}$, which consist of a corresponding LDR frame and its neighbors. Following the previous work [5], five frames are used as input in the case of having two alternating exposures ($N = 2$), and seven frames for the three ($N = 3$). For brevity, this paper mostly addresses the case of two alternating exposures and discusses the extension to the three in Section 3.4.

Before feeding the inputs into our framework, we map the LDR frames to the HDR using gamma correction with

$\gamma = 2.2$:

$$X_t = L_t^\gamma / e_t, \quad (1)$$

where e_t is the exposure time of L_t . We then construct a 6-channel input by concatenating L_t and X_t along the channel dimension. In the rest of the paper, we denote the LDR input as L_t for brevity, though its gamma-transformed HDR image X_t is actually used together as the input.

Our framework is composed of weight-sharing luminance-based alignment networks (LANs) and a merging network. Each LAN aligns the motion of the supporting frame L_{t+i} to that of the reference L_t , and the merging network combines the features from the LANs into an HDR frame H_t . As shown in Fig. 2, the LAN has two different modules. The alignment module calculates the attention score between the reference frame (*query*) and the supporting frame (*key*) and then aligns the supporting frame (*value*) based on this score. Since the alignment module can produce incomplete content in saturated areas, the hallucination module generates fine details in the desired regions. The features from the two modules are fused in an adaptive blending layer. The outputs of the LANs are concatenated and fed to the merging network for HDR frame composition. The details of each network are described in the following sections.

3.1. Luminance-based Alignment Network (LAN)

Alignment module The core idea of the alignment module is to calculate attention between a neighboring frame (*key*) and a reference frame (*query*) based on their content and then rearrange the neighboring frame (*value*) using the estimated attention. To alleviate memory usage and computation for attention operation, we input low-resolution images downsampled by a factor of 4. Here, we first adjust the exposure of the downsampled reference frame L_t^\downarrow to that of the downsampled neighboring frame L_{t+i}^\downarrow : $g_{t+i}(L_t^\downarrow) = \text{clip}\left(L_t^\downarrow (e_{t+i}/e_t)^{1/\gamma}\right)$, where e_t and e_{t+i} denote the exposure time of L_t^\downarrow and L_{t+i}^\downarrow , respectively. Then, after conversion to YCbCr space, only Y channels of each image, $L_{t+i}^{\downarrow,Y}$ and $g_{t+i}(L_t^\downarrow)^Y$, are fed to the key/query extractor F_{KQ} . The F_{KQ} can extract key and query features focusing on content rather than superficial color information.

The key and query features are then unfolded into patches to reflect neighbors. Likewise, the downsampled neighboring frame L_{t+i}^\downarrow is embedded through a value extractor F_V and unfolded to construct value features. The key K , query Q , and value V for attention operation can be represented as a set of patches: $K = \{\mathbf{k}_i, i \in [1, n]\}$, $Q = \{\mathbf{q}_i, i \in [1, n]\}$, and $V = \{\mathbf{v}_i, i \in [1, n]\}$, respectively. Each patch has the dimension of $3 \times 3 \times C$, and n is the number of pixels in the key/query/value features.

To estimate correlation between Q and K , patch-wise



Figure 3: An example of a blending map M when aligning a neighboring frame L_{t+1} to the reference L_t with saturation. The darker the map, the more features from the hallucination module F_h contribute to reconstruction.

cosine similarity is calculated as:

$$c_{i,j} = \left\langle \frac{\mathbf{q}_i}{\|\mathbf{q}_i\|}, \frac{\mathbf{k}_j}{\|\mathbf{k}_j\|} \right\rangle. \quad (2)$$

In order to prevent blurry results, we exploit the most correlated (top-1) element, constructing an index map $I = \{p_i, i \in [1, n]\}$ and a confidence map $S = \{s_i, i \in [1, n]\}$. For each query patch \mathbf{q}_i , an index of the most correlated key patch p_i and its confidence score s_i are determined as:

$$p_i = \underset{j}{\operatorname{argmax}} c_{i,j}, \quad (3)$$

$$s_i = \max_j c_{i,j}. \quad (4)$$

Then, rearranged value features V' are obtained, where v_{p_i} is located at the i -th location. To prevent information loss in the case of mismatching, we concatenate the original value features V with the rearranged features V' and then multiply the confidence map S . The adjusted features are up-sampled to the resolution of the input L_{t+i}^\downarrow progressively through upsampling blocks, resulting in the aligned features F_a .

Hallucination module It is difficult to perform the alignment when the reference image has severely saturated areas such as the one in Fig. 3. Also, downsampled inputs lack sharp details while improving efficiency. The hallucination module aims to fill in the missing content resulting from saturation and generate high-frequency details. To this end, we utilize gated convolution [50], which is widely adopted in inpainting networks, to restore specific saturated regions. Also, masks indicating saturated parts as well as the original images are used as input to an encoder-decoder architecture. According to the brightness information of the input masks, the gated convolution performs adaptively with respect to spatial and channel dimensions. Given input features F_i , the output of the gated convolution F_o is formulated as:

$$F_o = \rho\left(\sum \sum W_f \cdot F_i\right) \odot \sigma\left(\sum \sum W_g \cdot F_i\right), \quad (5)$$

where \odot denotes element-wise multiplication. ρ and σ are the exponential linear unit and sigmoid, respectively. W_f denotes a convolutional filter for the feature, and W_g is the

one for gating. Thus, gated convolution enables dynamic operation that is aware of saturation. We use the Y channel of each image as a mask to provide explicit but continuous information about luminance. The luminance channel indicates dark regions as well as saturated parts, thus helping to suppress noise in dark areas. Analysis of the input to the hallucination module is presented in Section 4.3.

Adaptive blending layer Aligned features from the alignment module F_a are fused with intermediate features of the hallucination module F_h through an adaptive blending layer. For the two modules to complement each other, the adaptive blending layer estimates a blending map $M \in \mathbb{R}^{H \times W \times 1}$ in the range $[0, 1]$, where $H \times W$ is the spatial resolution of F_a and F_h . The output features of the fusion F_{out} are represented as follows:

$$F_{out} = (1 - M) \odot F_h + M \odot (F_h + F_a). \quad (6)$$

When the alignment is challenging in saturated parts due to missing information, the contribution of F_a becomes small, represented as a low value in the map M . (See an example in Fig. 3.) The F_{out} passes through the remaining upsampling blocks to recover its spatial size.

3.2. Merging Network

Since the proposed LANs output well-aligned features, our merging network has a simple structure. The aligned features are concatenated along the channel axis and fed to a convolutional layer and five Res FFT-Conv Blocks [31]. This network can cover the entire frame globally by performing convolutions in the frequency domain. After the final convolutional layer with a sigmoid activation function, the final HDR frame H_t is constructed.

3.3. Training Loss

Since HDR images are usually displayed after tonemapping, we calculate loss functions between the tonemapped predicted HDR frame $\mathcal{T}(\hat{H}_t)$ and the tonemapped ground truth HDR image $\mathcal{T}(H_t)$ using the differentiable μ -law:

$$\mathcal{T}(H) = \frac{\log(1 + \mu H)}{\log(1 + H)}, \quad (7)$$

where μ is set to 5000.

\mathcal{L}_1 loss We adopt \mathcal{L}_1 loss which is defined as: $\mathcal{L}_1 = \|\mathcal{T}(\hat{H}_t) - \mathcal{T}(H_t)\|_1$.

Perceptual loss Using the feature map from the pre-trained VGG-19 network $\phi(\cdot)$, we calculate perceptual loss \mathcal{L}_{per} as: $\mathcal{L}_{per} = \|\phi(\mathcal{T}(\hat{H}_t)) - \phi(\mathcal{T}(H_t))\|_2^2$.

Frequency loss We also utilize a frequency loss which has been shown to be effective in recent low-level vision tasks [12, 6, 43]. The frequency loss \mathcal{L}_{freq} is defined as:

$\mathcal{L}_{freq} = \|\mathcal{F}(\mathcal{T}(\hat{H}_t)) - \mathcal{F}(\mathcal{T}(H_t))\|_1$, where \mathcal{F} is the fast Fourier transform (FFT).

Temporal loss To generate a perceptually natural video, temporal consistency without abrupt change is crucial. Since every HDR frame is constructed independently, we present our temporal loss \mathcal{L}_{temp} to improve temporal consistency. Specifically, we enforce a difference map between the consecutive tonemapped outputs to be similar to the ground truth using the Charbonnier penalty function [4]. The temporal loss for this enforcement is defined as:

$$\mathcal{L}_{temp} = \sqrt{\|(\mathcal{T}(\hat{H}_t) - \mathcal{T}(\hat{H}_{t-1})) - (\mathcal{T}(H_t) - \mathcal{T}(H_{t-1}))\|_2^2 + \epsilon^2}, \quad (8)$$

where ϵ is set as 10^{-3} .

Our total loss is represented as: $\mathcal{L}_{total} = \lambda_1 \mathcal{L}_1 + \lambda_{per} \mathcal{L}_{per} + \lambda_{freq} \mathcal{L}_{freq} + \lambda_{temp} \mathcal{L}_{temp}$, where $\lambda_1 = 1$, $\lambda_{per} = 0.1$, $\lambda_{freq} = 0.1$, and $\lambda_{temp} = 0.1$.

3.4. Extension to Three Exposures

Given an LDR sequence with alternating exposures in three levels, seven LDR frames are used to compose an HDR frame. Our LAN aligns each pair of a supporting frame and the reference, and the merging network combines them in the same manner.

4. Experiments

4.1. Experimental Setup

Datasets We use synthetic training data constructed using the Vimeo-90K septuplet dataset [46]. Since the Vimeo-90K dataset is not optimized for HDR video reconstruction, we convert the original data to LDR sequences with alternating exposures following the previous work [5]. We evaluate our framework on two synthetic videos (POKER FULLSHOT and CAROUSEL FIREWORKS) from the Cinematic Video dataset [11] and DeepHDRVideo dataset [5]. The DeepHDRVideo dataset [5] includes both real-world dynamic scenes and static scenes augmented with random global motion. HDRVideo dataset [20] is used only for qualitative evaluation due to the absence of ground truth.

Training details We adopt AdamW optimizer [27] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$, and the learning rate is set to 10^{-4} . We augment training patches using random flipping, rotation by multiples of 90° , and color augmentation. More details are included in the supplementary material.

Evaluation metrics We compute PSNR_T , SSIM_T , PSNR_{PU} , SSIM_{PU} , and HDR-VDP-2 between the predicted results and ground truth frames. PSNR_T and SSIM_T are computed on the tonemapped images using the μ -law. PSNR_{PU} and SSIM_{PU} are calculated after perceptually uniform encoding [30], which maps the peak value in the HDR image to the peak HDR display luminance

Table 1: Quantitative comparisons of our method with other state-of-the-arts on the Cinematic Video dataset [11].

Methods	2 Exposures					3 Exposures				
	PSNR _T	SSIM _T	PSNR _{PU}	SSIM _{PU}	HDR-VDP-2	PSNR _T	SSIM _T	PSNR _{PU}	SSIM _{PU}	HDR-VDP-2
Kalantari [20]	37.51	0.9016	39.13	0.9319	60.16	30.36	0.8133	32.53	0.8251	57.68
Kalantari [19]	37.06	0.9053	39.10	0.9286	70.82	33.21	0.8402	35.38	0.8582	62.44
Yan [47]	31.65	0.8757	32.99	0.7523	69.05	34.22	0.8604	36.04	0.8357	66.18
Prabhakar [36]	34.72	0.8761	36.06	0.8257	68.82	34.02	0.8633	36.19	0.8590	65.00
Chen [5]	35.65	0.8949	37.12	0.8156	72.09	34.15	0.8847	36.23	0.8357	66.81
Ours	38.22	0.9100	40.04	0.9039	69.15	35.07	0.8695	37.22	0.8666	65.42

Table 2: Quantitative comparisons of our method with other state-of-the-arts on the DeepHDRVideo dataset [5].

Methods	2 Exposures					3 Exposures				
	PSNR _T	SSIM _T	PSNR _{PU}	SSIM _{PU}	HDR-VDP-2	PSNR _T	SSIM _T	PSNR _{PU}	SSIM _{PU}	HDR-VDP-2
Kalantari [20]	40.33	0.9409	43.71	0.9646	66.11	38.45	0.9489	42.35	0.9740	57.31
Kalantari [19]	39.91	0.9329	43.31	0.9641	71.11	38.78	0.9331	41.80	0.9647	65.73
Yan [47]	40.54	0.9452	45.33	0.9616	69.67	40.20	0.9531	42.36	0.9700	68.23
Prabhakar [36]	40.21	0.9414	45.16	0.9593	70.27	39.48	0.9453	41.15	0.9666	65.93
Chen [5]	42.48	0.9620	45.79	0.9773	74.80	39.44	0.9569	41.57	0.9725	67.76
Ours	41.59	0.9472	44.43	0.9730	71.34	40.48	0.9504	42.38	0.9755	68.61

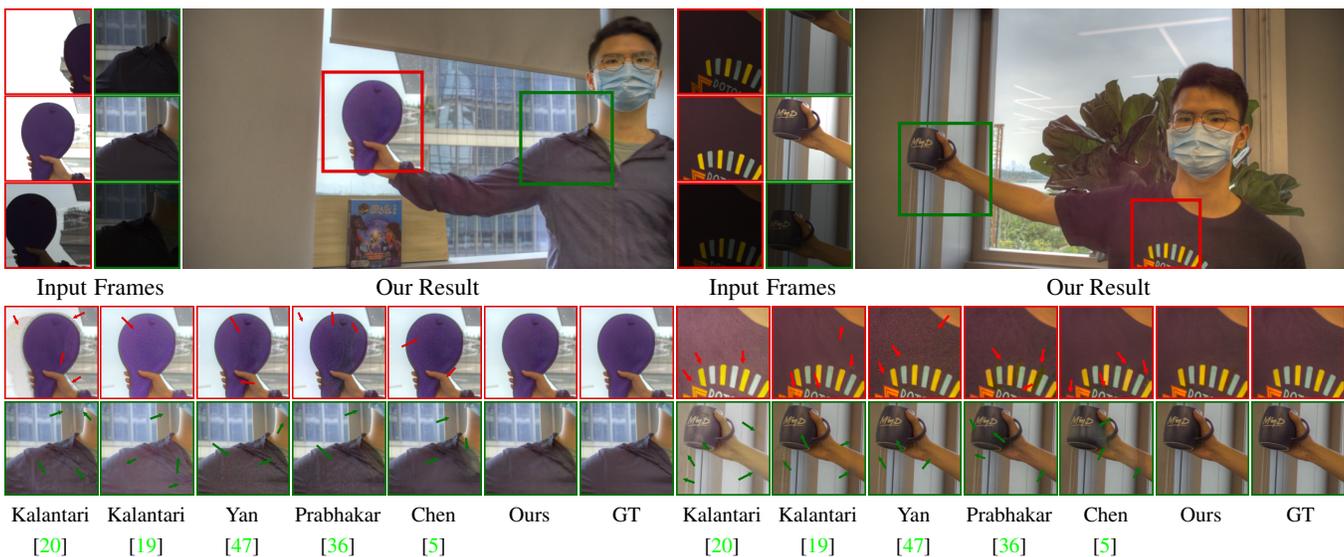


Figure 4: Qualitative results on 3-exposure sequences in the DeepHDRVideo dataset [5]. Only the middle three input frames are displayed.

$4000\text{cd}/\text{m}^2$. When computing the HDR-VDP-2, the angular resolution of the image in terms of the number of pixels per visual degree is set as 30.

4.2. Comparisons

We compare our results with previous HDR video reconstruction approaches [20, 19, 5] and state-of-the-art HDR image composition methods [47, 36] after adapting them to the video reconstruction task. We use the official codes if available, otherwise, we re-implement them based on the papers. For a fair comparison, we have trained deep learning-based methods with the dataset we use. All the visual results are tonemapped using Photomatrix [34] and best

viewed by zooming into the electronic version.

Quantitative evaluations Our results on Cinematic Video dataset [11] and DeepHDRVideo dataset [5] are listed in Table 1 and Table 2, respectively, where our approach achieves superior or comparable performances to state-of-the-art methods. Especially, in the case of data with 3 exposures, our method consistently demonstrates precise reconstruction ability by handling large variations of motions and brightness.

Qualitative evaluations We visualize our results on sequences with saturation and large motions. Fig. 4 shows the results on two 3-exposure sequences from the DeepHDRVideo dataset [5]. The other methods except ours fail

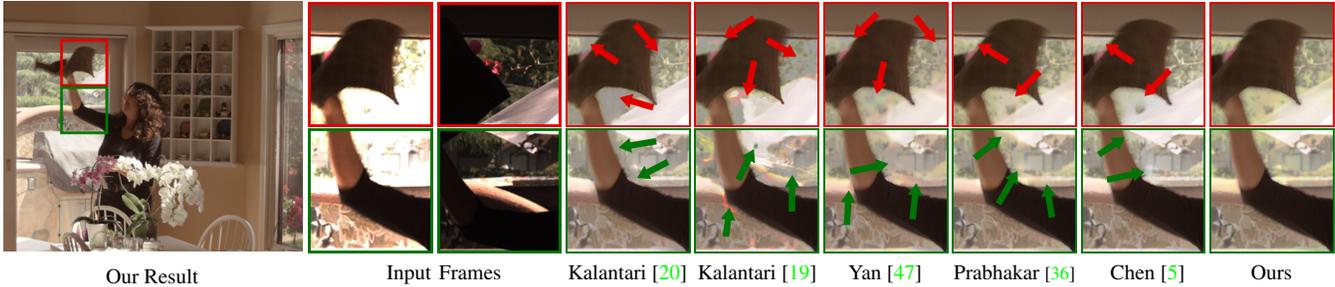


Figure 5: Qualitative comparisons on the sequence with 2 exposures from the HDRVideo dataset [20]. Only the middle two input frames are displayed.

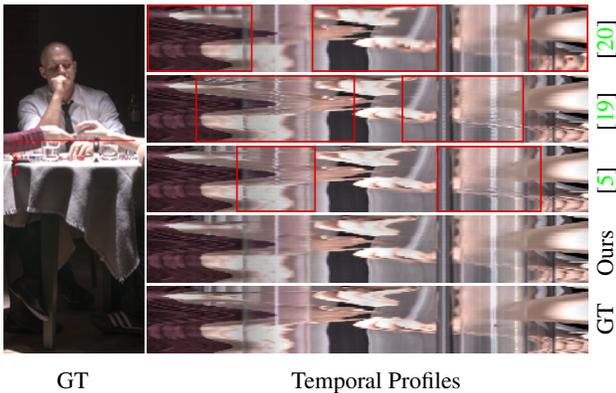


Figure 6: Visualization of temporal consistency.

Table 3: Comparisons of inference time.

	2 Exposures		3 Exposures	
	1280 × 720	1536 × 813	1280 × 720	1536 × 813
Kalantari [20]	140.02s	166.10s	321.16s	414.44s
Kalantari [19]	0.37s	0.69s	0.49s	0.72s
Chen [5]	0.54s	1.25s	0.54s	1.30s
Ours	0.37s	0.61s	0.57s	0.82s

to avoid ghosting artifacts and color distortions when fast motion occurs with exposure changes. Especially, the approaches using flow-based estimation [19, 5] suffer from artifacts caused by flow estimation error. We also evaluate on a sequence with two alternating exposures from the HDRVideo dataset [20]. Fig. 5 shows that our method generates fine details even when large occlusion and saturation exist. More visual results are presented in Fig. 1 and the supplementary material.

Temporal profiles This section provides temporal profiles to demonstrate the temporal consistency of our results. We produce temporal profiles by recording a single line across all the frames and stitching them in order. Fig. 6 shows the results on a 3-exposure sequence in the Cinematic Video dataset [11]. Our method generates temporally consistent videos with minimal flickering artifacts and preserves detailed texture successfully.

Table 4: Effectiveness of each component in the LAN. AM, HM, and ABL denote the alignment module, the hallucination module, and the adaptive blending layer, respectively.

Model	AM	HM	ABL	PSNR _T	PSNR _{PU}
(a)	✓			36.37	37.99
(b)		✓		36.32	36.68
(c)	✓	✓		37.57	38.96
(d)	✓	✓	✓	38.22	40.04

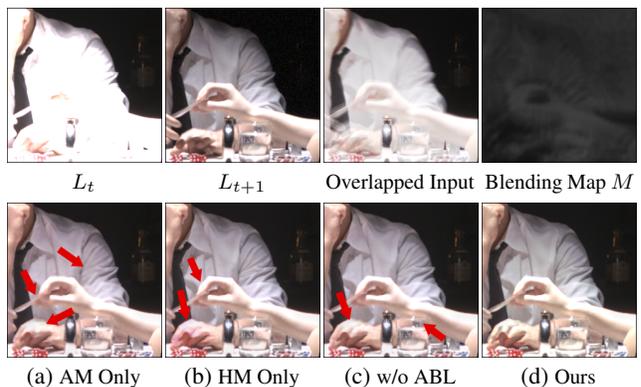


Figure 7: Analysis of the components in the LAN.

Computation time We compare the inference time of our method with other HDR video construction methods in Table 3. We use a single NVIDIA TITAN Xp GPU, while the optimization-based method of Kalantari *et al.* [20] is run on CPUs. Our approach is faster than the method of Chen *et al.* [5], and the gap increases as the image resolution becomes larger.

4.3. Analysis

In this section, we analyze the effects of the proposed components with the example of sequences having two alternating exposures. All the quantitative evaluations are conducted on the Cinematic Video dataset [11].

LAN Our LAN is composed of the alignment module (AM) and the hallucination module (HM), and features from

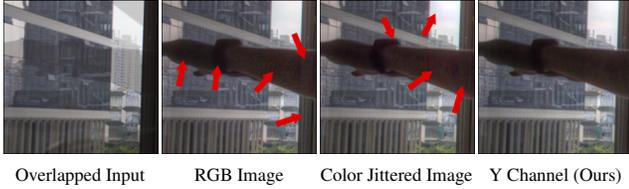


Figure 8: Results with different inputs to the key/query extractor of the alignment module.

Table 5: Effectiveness of gated convolution and input masks in the hallucination module.

Model	Convolution	Input Mask	PSNR _T	PSNR _{PU}
(A)	Regular	Y channel	37.02	38.55
(B)	Gated	-	36.49	38.10
(C)	Gated	Binary	36.88	38.32
(D)	Gated	Y channel	38.22	40.04

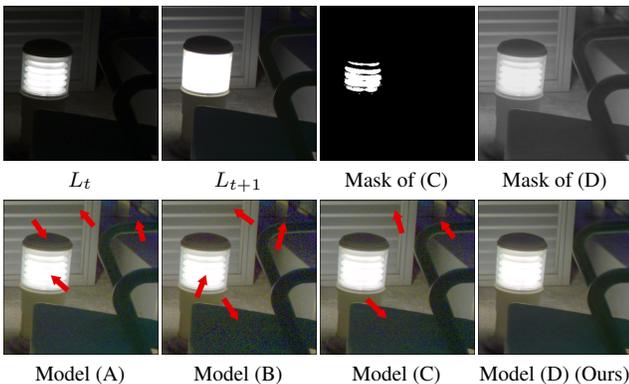


Figure 9: Results with different convolutions and input masks in the hallucination module.

the two modules are fused in the adaptive blending layer (ABL). As shown in Table 4, the model with only AM or HM achieves far lower PSNRs than the full proposed model. The model which has both AM and HM, but blends two features using a simple concatenation, improves performance. But it suffers from undesirable artifacts in a challenging case such as a sequence with alternating exposures whose middle frame has large saturations as in Fig. 7. In this case, the AM cannot perform a perfect alignment because a large amount of content in the reference frame is totally washed out and thus produces blurry texture, as shown in Fig. 7a. On the other hand, Fig. 7b shows that the HM generates sharp details but fails to align the motions according to the reference. (See the hand holding a card.) Our full model, including the AM, HM, and ABL (Fig. 7d), outputs better results than the individual modules by properly combining two kinds of features. For example, our blending map has low values around the moving object in saturated parts, demonstrating that inaccurate AM features are suppressed in those areas.

Table 6: Effectiveness of the proposed temporal loss.

Methods	PSNR _T	SSIM _T	PSNR _{PU}	SSIM _{PU}
w/o \mathcal{L}_{temp}	36.33	0.9200	37.99	0.8972
w/ \mathcal{L}_{temp}	38.22	0.9100	40.04	0.9039

Inputs for alignment module In our AM, we perform content-based alignment using Y channels of two input frames to prevent a simplistic matching based on color information. We compare our model with the one that uses original RGB images for alignment in Fig. 8. When using RGB images for attention operation, registration is unsuccessful when objects with similar colors overlap. We also conduct experiments with color-jittered RGB images as another means of relying less on color and more on content. The model using the color-jittered images improves alignment, but our model shows the best alignment results.

Hallucination module Our HM in the LAN utilizes gated convolution and input masks to operate adaptively to luminance. The effects of these elements are shown in Table 5 and Fig. 9. Model (A) with regular convolution takes the Y channel as input, but the dynamic operation is not possible. Model (B) includes gated convolution, but there is no explicit instruction for adaptive convolution, resulting in unsatisfying results. For model (C), a binary mask obtained from thresholding at a predefined value of 0.95 is given. The binary mask indicates saturated parts but provides limited information regarding brightness. Fig. 9 shows a scene with both highlighted and very dark regions. It is noteworthy that other models not only lack details in bright parts but also generate significant noise in dark areas. In contrast, our framework suppresses noise in dark regions and produces detailed content in saturated regions adaptively.

Temporal loss We exploit the temporal loss \mathcal{L}_{temp} to generate a temporally coherent HDR video. The temporal loss forces difference of consecutive frames to be similar to one of the ground truth. Table 6 shows that the overall quantitative performance is improved by using the temporal loss.

5. Conclusion

We have proposed an end-to-end HDR video composition framework using LDR videos with alternating exposures. To enhance alignment accuracy and restore washed-out regions, our alignment network performs content-based alignment with attention and recovers details in saturated areas using brightness-adaptive convolution. The aligned features and the hallucinated features are fused adaptively using a learned blending map. The outputs of the alignment module are reconstructed into an HDR frame through the merging network. During training, the temporal loss is adopted for temporally consistent results. The overall results demonstrate that the proposed method generates high-quality HDR videos.

Acknowledgments

This research was supported in part by Samsung Electronics Co., Ltd., in part by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (2021R1A2C2007220), and partially by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2021-0-01062).

References

- [1] Masheal M Alghamdi, Qiang Fu, Ali Kassem Thabet, and Wolfgang Heidrich. Reconfigurable snapshot hdr imaging using coded masks and inception network. 2019. **1**
- [2] Jaehyun An, Seong Jong Ha, and Nam Ik Cho. Probabilistic motion pixel detection for the reduction of ghost artifacts in high dynamic range images from multiple exposures. *EURASIP Journal on Image and Video Processing*, 2014(1):1–15, 2014. **1, 2**
- [3] Francesco Banterle, Patrick Ledda, Kurt Debattista, and Alan Chalmers. Inverse tone mapping. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia*, pages 349–356, 2006. **1**
- [4] Pierre Charbonnier, Laure Blanc-Feraud, Gilles Aubert, and Michel Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of 1st International Conference on Image Processing*, volume 2, pages 168–172. IEEE, 1994. **5**
- [5] Guanying Chen, Chaofeng Chen, Shi Guo, Zhetong Liang, Kwan-Yee K Wong, and Lei Zhang. Hdr video reconstruction: A coarse-to-fine network and a real-world benchmark dataset. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2502–2511, 2021. **1, 2, 3, 5, 6, 7**
- [6] Sung-Jin Cho, Seo-Won Ji, Jun-Pyo Hong, Seung-Won Jung, and Sung-Jea Ko. Rethinking coarse-to-fine approach in single image deblurring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4641–4650, 2021. **5**
- [7] Ugur Cogalan, Mojtaba Bemana, Karol Myszkowski, Hans-Peter Seidel, and Tobias Ritschel. Learning hdr video reconstruction for dual-exposure sensors with temporally-alternating exposures. *Computers & Graphics*, 105:57–72, 2022. **3**
- [8] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 369–378. ACM Press/Addison-Wesley Publishing Co., 1997. **1**
- [9] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafal K Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *ACM Transactions on Graphics (TOG)*, 36(6):178, 2017. **1**
- [10] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Transactions on Graphics (Proc. of SIGGRAPH ASIA 2017)*, 36(6), Nov. 2017. **1**
- [11] Jan Froehlich, Stefan Grandinetti, Bernd Eberhardt, Simon Walter, Andreas Schilling, and Harald Brendel. Creating cinematic wide gamut hdr-video for the evaluation of tone mapping operators and hdr-displays. In *Digital photography X*, volume 9023, pages 279–288. SPIE, 2014. **5, 6, 7**
- [12] Dario Fuoli, Luc Van Gool, and Radu Timofte. Fourier space losses for efficient perceptual image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2360–2369, 2021. **5**
- [13] Thorsten Grosch et al. Fast and robust high dynamic range image generation with camera and object movement. *Vision, Modeling and Visualization, RWTH Aachen*, 277284, 2006. **2**
- [14] Yulia Gryaditskaya, Tania Pouli, Erik Reinhard, Karol Myszkowski, and Hans-Peter Seidel. Motion aware exposure bracketing for hdr video. In *Computer Graphics Forum*, volume 34, pages 119–130. Wiley Online Library, 2015. **3**
- [15] Yong Seok Heo, Kyoung Mu Lee, Sang Uk Lee, Youngsu Moon, and Joonhyuk Cha. Ghost-free high dynamic range imaging. In *Asian Conference on Computer Vision*, pages 486–500. Springer, 2010. **2**
- [16] Jun Hu, Orazio Gallo, Kari Pulli, and Xiaobai Sun. Hdr deghosting: How to deal with saturation? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1163–1170, 2013. **2**
- [17] Yitong Jiang, Inchang Choi, Jun Jiang, and Jinwei Gu. Hdr video reconstruction with tri-exposure quad-bayer sensors. *arXiv preprint arXiv:2103.10982*, 2021. **3**
- [18] Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *ACM Trans. Graph.*, 36(4):144–1, 2017. **1, 2, 3**
- [19] Nima Khademi Kalantari and Ravi Ramamoorthi. Deep hdr video from sequences with alternating exposures. In *Computer graphics forum*, volume 38, pages 193–205. Wiley Online Library, 2019. **2, 3, 6, 7**
- [20] Nima Khademi Kalantari, Eli Shechtman, Connelly Barnes, Soheil Darabi, Dan B Goldman, and Pradeep Sen. Patch-based high dynamic range video. *ACM Trans. Graph.*, 32(6):202–1, 2013. **2, 3, 5, 6, 7**
- [21] Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High dynamic range video. In *ACM Transactions on Graphics (TOG)*, volume 22, pages 319–325. ACM, 2003. **1, 2, 3**
- [22] Erum Arif Khan, Ahmet Oguz Akyuz, and Erik Reinhard. Ghost removal in high dynamic range images. In *2006 International Conference on Image Processing*, pages 2005–2008. IEEE, 2006. **2**
- [23] Joel Kronander, Stefan Gustavson, Gerhard Bonnet, Anders Ynnerman, and Jonas Unger. A unified framework for multi-sensor hdr video reconstruction. *Signal Processing: Image Communication*, 29(2):203–215, 2014. **1, 3**
- [24] Chul Lee, Yuelong Li, and Vishal Monga. Ghost-free high dynamic range imaging via rank minimization. *IEEE Signal Processing Letters*, 21(9):1045–1049, 2014. **2**
- [25] Yuelong Li, Chul Lee, and Vishal Monga. A maximum a posteriori estimation framework for robust high dynamic range video synthesis. *IEEE Transactions on Image Processing*, 26(3):1143–1157, 2016. **3**

- [26] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image hdr reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1651–1660, 2020. **1**
- [27] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. **5**
- [28] Stephen Mangiat and Jerry Gibson. High dynamic range video with ghost removal. In *Applications of Digital Image Processing XXXIII*, volume 7798, pages 307–314. SPIE, 2010. **2, 3**
- [29] Stephen Mangiat and Jerry Gibson. Spatially adaptive filtering for registration artifact removal in hdr video. In *2011 18th IEEE International Conference on Image Processing*, pages 1317–1320. IEEE, 2011. **2, 3**
- [30] Rafal Mantiuk and Maryam Azimi. Pu21: A novel perceptually uniform encoding for adapting existing quality metrics for hdr. 2021. **5**
- [31] Xintian Mao, Yiming Liu, Wei Shen, Qingli Li, and Yan Wang. Deep residual fourier transformation for single image deblurring. *arXiv preprint arXiv:2111.11745*, 2021. **5**
- [32] Demetris Marnerides, Thomas Bashford-Rogers, Jonathan Hatchett, and Kurt Debattista. Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. In *Computer Graphics Forum*, volume 37, pages 37–49. Wiley Online Library, 2018. **1**
- [33] Tae-Hyun Oh, Joon-Young Lee, Yu-Wing Tai, and In So Kweon. Robust high dynamic range imaging by rank minimization. *IEEE transactions on pattern analysis and machine intelligence*, 37(6):1219–1232, 2015. **1, 2**
- [34] Photomatix. Commercially-available hdr processing software. <http://www.hdrsoft.com>, 2017. **6**
- [35] K Ram Prabhakar, Susmit Agrawal, Durgesh Kumar Singh, Balraj Ashwath, and R Venkatesh Babu. Towards practical and efficient high-resolution hdr deghosting with cnn. In *European Conference on Computer Vision*, pages 497–513. Springer, 2020. **1**
- [36] K Ram Prabhakar, Gowtham Senthil, Susmit Agrawal, R Venkatesh Babu, and Rama Krishna Sai S Gorthi. Labeled from unlabeled: Exploiting unlabeled data for few-shot deep hdr deghosting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4875–4885, 2021. **2, 6, 7**
- [37] Zhiyuan Pu, Peiyao Guo, M Salman Asif, and Zhan Ma. Robust high dynamic range (hdr) imaging with complex motion and parallax. In *Proceedings of the Asian Conference on Computer Vision*, 2020. **2**
- [38] Allan G Rempel, Matthew Trentacoste, Helge Seetzen, H David Young, Wolfgang Heidrich, Lorne Whitehead, and Greg Ward. Ldr2hdr: on-the-fly reverse tone mapping of legacy video and photographs. *ACM transactions on graphics (TOG)*, 26(3):39–es, 2007. **1**
- [39] Pradeep Sen, Nima Khademi Kalantari, Maziar Yaesoubi, Soheil Darabi, Dan B Goldman, and Eli Shechtman. Robust patch-based hdr reconstruction of dynamic scenes. *ACM Trans. Graph.*, 31(6):203–1, 2012. **2, 3**
- [40] Michael D Tocci, Chris Kiser, Nora Tocci, and Pradeep Sen. A versatile hdr video production system. *ACM Transactions on Graphics (TOG)*, 30(4):1–10, 2011. **1, 3**
- [41] Jonas Unger and Stefan Gustavson. High-dynamic-range video for photometric measurement of illumination. In *Sensors, Cameras, and Systems for Scientific/Industrial Applications VIII*, volume 6501, pages 106–115. SPIE, 2007. **1, 3**
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. **2**
- [43] Jiaming Wang, Zhenfeng Shao, Xiao Huang, Tao Lu, Ruiqian Zhang, and Yong Li. From artifact removal to super-resolution. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–15, 2022. **5**
- [44] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. **3**
- [45] 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)*, pages 117–132, 2018. **1, 2**
- [46] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019. **5**
- [47] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 1751–1760, 2019. **1, 2, 6, 7**
- [48] Qingsen Yan, Lei Zhang, Yu Liu, Yu Zhu, Jinqiu Sun, Qinfeng Shi, and Yanning Zhang. Deep hdr imaging via a non-local network. *IEEE Transactions on Image Processing*, 29:4308–4322, 2020. **2**
- [49] Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, and Rynson W.H. Lau. Image correction via deep reciprocating hdr transformation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. **1**
- [50] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4471–4480, 2019. **2, 4**
- [51] Wei Zhang and Wai-Kuen Cham. Gradient-directed multiexposure composition. *IEEE Transactions on Image Processing*, 21(4):2318–2323, 2012. **2**
- [52] Henning Zimmer, Andrés Bruhn, and Joachim Weickert. Freehand hdr imaging of moving scenes with simultaneous resolution enhancement. In *Computer Graphics Forum*, volume 30, pages 405–414. Wiley Online Library, 2011. **2**