

HDR-GAN: HDR Image Reconstruction from Multi-Exposed LDR Images with Large Motions

Yuzhen Niu, Jianbin Wu, Wenxi Liu, Wenzhong Guo, Rynson W.H. Lau

Abstract—Synthesizing high dynamic range (HDR) images from multiple low-dynamic range (LDR) exposures in dynamic scenes is challenging. There are two major problems caused by the large motions of foreground objects. One is the severe misalignment among the LDR images. The other is the missing content due to the over-/under-saturated regions caused by the moving objects, which may not be easily compensated for by the multiple LDR exposures. Thus, it requires the HDR generation model to be able to properly fuse the LDR images and restore the missing details without introducing artifacts. To address these two problems, we propose in this paper a novel GAN-based model, *HDR-GAN*, for synthesizing HDR images from multi-exposed LDR images. To our best knowledge, this work is the first GAN-based approach for fusing multi-exposed LDR images for HDR reconstruction. By incorporating adversarial learning, our method is able to produce faithful information in the regions with missing content. In addition, we also propose a novel generator network, with a reference-based residual merging block for aligning large object motions in the feature domain, and a deep HDR supervision scheme for eliminating artifacts of the reconstructed HDR images. Experimental results demonstrate that our model achieves state-of-the-art reconstruction performance over the prior HDR methods on diverse scenes.

Index Terms—High dynamic range imaging, generative adversarial networks, multi-exposed imaging

I. INTRODUCTION

Although the human visual system has a much higher visual dynamic range, most off-the-shelf digital cameras typically produce photos with a limited range of illumination, which may not be desirable in many scenarios. Some specialized hardware devices [1], [2] have been proposed to produce high dynamic range (HDR) images directly, but they are usually too expensive to be widely used. In recent years, with prevailing mobile devices, there is a high demand for capturing HDR images of scenes using light-weighted and low-cost monocular mobile cameras, in order to produce photos that cover a broad range of illumination.

One popular approach to obtain an HDR image is to merge several low dynamic range (LDR) images captured from multiple exposures [3]–[6]. Given a set of multi-exposed LDR images, one of them (typically with the medium exposure) is used as the reference image and the rest are used to compensate for the missing details due to over-/under-exposure of some local regions. When the LDR images are perfectly aligned pixel-wisely, these traditional methods that rely on

hand-crafted features can produce high quality HDR images. However, foreground and background misalignments are difficult to avoid in practice due to object motion, causing the output HDR images to have blur and ghosting artifacts [7], [8]. In particular, there are two main problems caused by the large motions of foreground objects. One is the severe misalignment among the LDR images, and the other is the difficulty in compensating the missing content in the over-/under-exposed regions via multiple exposures due to occlusions of the moving objects in the images, as shown in Fig. 1.

To address these problems, some traditional methods attempt to align the LDR images in the preprocessing stage [9], [10] before fusing them, e.g., using optical flow [7], [10]. With the popularity of deep learning, CNN-based HDR methods are proposed [11]–[14]. These methods have strong ability in hallucinating details for the reconstructed HDR images. However, their network structures do not specifically handle the misalignment of the LDR images. For example, Kalantari *et al.* [11] directly utilize a CNN for pixel-wise merging. Wu *et al.* [12] apply the U-Net or ResNet architecture with sparse skip connections for LDR fusion. Yan *et al.* [13] introduces dense connections into the model, without considering multi-scale context information of the LDR images. These methods cannot guarantee that the fused features of the LDR images are aligned well. Thus, they still require to incorporate homography transformation or optical flow to holistically align the LDR images beforehand, to cover the shortage of their networks. However, when the LDR images contain large motions or significant misalignment, they may still suffer from having artifacts due to the unreliability of optical flow estimation, especially for images captured with different exposure levels, as demonstrated in Fig. 1.

In this work, we address the two aforementioned problems by proposing a novel GAN-based model, *HDR-GAN*, which produces high-quality HDR images from multi-exposed LDR images without the need to explicitly align the LDR images. To our best knowledge, our approach is the first attempt to adapt a GAN-based model in this task for fusing multi-exposed LDR images to reconstruct an HDR image, which can help recover regions with incomplete image content due to under-/over-exposures. To address the limitations of prior CNN-based methods in fusing the misaligned LDR features, in our GAN-based framework, we propose a reference-based residual merging module for aligning large object motions in the feature domain and a deep HDR supervision scheme for eliminating artifacts in the reconstructed HDR image. In particular, the proposed reference-based residual merging module contains reference-based residual structures to trans-

Yuzhen Niu, Jianbin Wu, Wenzhong Guo, and Wenxi Liu are with the College of Mathematics and Computer Science, Fuzhou University, China. Wenxi Liu is the corresponding author.

Rynson W.H. Lau is with the Department of Computer Science, City University of Hong Kong.

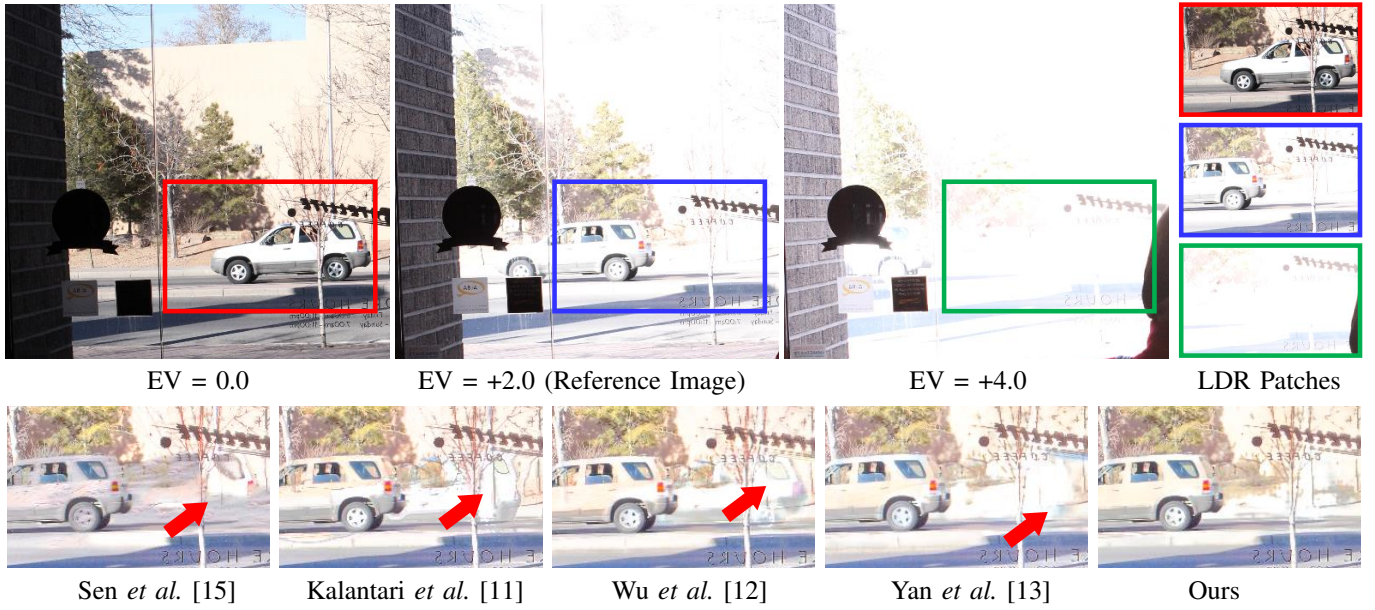


Fig. 1. Our GAN-based model is able to fuse multiple LDR images with large object motions into a ghost-free HDR image, without the need to explicitly align the images. In this example, a car is moving rapidly from right to left. In the over-exposure of the third shot (i.e., the green patch), the region where the car used to be in the first shot (i.e., the red patch) is dis-occluded. Compared with the patch-based method [15] and recent deep learning based methods [11]–[13], which all suffer from various degrees of ghosting artifacts (highlighted by the red arrows), our approach is able to restore the missing details and produce a high-quality HDR image.

form the LDR features of multiple exposures to compensate for the reference image and thus implicitly align the features of low and high exposures to that of the median exposure (i.e., the reference image). In addition, inspired by the deep supervision of prior models [16]–[18], we introduce a deep HDR supervision scheme into our model. We first upsample the merged features of each scale, and sequentially pass them into the decoding block of its upper scale until reaching the top scale to generate several high-resolution HDR images for supervision. Supervising the resulting HDR images can then encourage the feature fusion of different scales, while eliminating the artifacts caused by the misalignment of the LDR image features.

To further encourage high-quality HDR reconstruction under challenging scenarios with severe missing details, we leverage the adversarial loss from [19] for discriminating the generated samples and the ground-truth by projecting them to a hypersphere space, which helps boost the image generation capability of our model. As demonstrated in Fig. 1, our model is able to hallucinate more details for the unobserved regions in the reference LDR image. To evaluate the performance of our model, we compare our approach with other HDR methods on a public benchmark and show that our model can achieve the state-of-the-art performance.

To sum up, the main contributions of our work include:

- We propose the first GAN-based method, *HDR-GAN*, for HDR reconstruction from multi-exposed LDR images. By incorporating adversarial learning, our method is able to produce faithful information for the HDR images when LDR images contain large object motions.
- We propose a novel generator network, with a reference-based residual merging block to implicitly align large

object/camera motions in the feature domain, and a deep HDR supervision scheme for eliminating artifacts.

- Our experimental results on the public benchmark demonstrate that the proposed model outperforms the state-of-the-art HDR models.

II. RELATED WORK

In this section, we briefly summarize related work on HDR reconstruction for static and dynamic scenes.

A. HDR Reconstruction for Static Scenes

There are HDR reconstruction methods that assume the multi-exposed input LDR images are aligned. These methods are essentially handling a static scene, in which all objects do not change during the capturing period. They typically use hand-crafted features or deep features to merge the LDR images pixel-wisely [20]–[25].

In recent years, with the progress of deep learning techniques, researchers attempt to address a more challenging problem of reconstructing an HDR image from a single LDR image [26]–[29]. For example, Eilertsen *et al.* [26] utilize the U-Net architecture to learn the mapping from LDR images to HDR images. Lee *et al.* [27] propose to generate LDR images with multiple exposures from a single LDR image using a GAN, and then merge the pseudo multi-exposed images pixel-wisely to produce an HDR image. Liu *et al.* [29] propose an HDR-to-LDR image formation pipeline that includes dynamic range clipping, non-linear mapping from a camera response function, and quantization. They then propose to learn three specialized CNNs to reverse these steps in order to reconstruct HDR images. Since it often exists missing details from a single image caused by quantization and saturation of the

camera sensor, the performance of these single image based HDR reconstruction methods generally do not perform well, compared with methods based on multiple exposures.

B. HDR Reconstruction for Dynamic Scenes

While LDR images with multiple exposures contain rich information of the scene, it is difficult to avoid object motions during the image capture period. A number of methods have been proposed to address this problem [7], [9]–[15], [30], [31]. Traditional HDR reconstruction methods for dynamic scenes estimate homography transformations to handle camera motions or to use optical flows as prior to align LDR images in the preprocessing stage [7], [10], [11].

Recent CNN-based methods formulate HDR reconstruction as an image translation problem from the LDR images to the HDR image [12]–[14], [30]. Wu *et al.* [12] first estimate the homography transformation to align the background of the LDR images, and then propose an U-Net or ResNet structure to directly learn the mapping from LDR images to a high-quality HDR image. Yan *et al.* [13] introduce spatial attention during the process of fusing LDR image features, while Yan *et al.* [14] leverage non-local blocks to process the global context of the unaligned image features. However, there are two challenging issues caused by the large motions of foreground objects. One is the severe misalignment of the LDR images, while the other is that the missing contents in the over-/under-exposed regions may not be easily compensated using multiple exposures, caused by the dis-occlusion of the moving objects in the scene. These existing methods either cannot guarantee to align the LDR images or fail to produce adequately faithful information for the missing image contents. Recently, there is a concurrent work [14] that explicitly addresses the feature alignment problem using a non-local network. However, due to the limited generation capability of the network structure, it is unable to hallucinate the missing contents caused by large object motions well.

In this work, we propose a GAN-based model to address the limitations of existing methods in handling large object motions in the scene.

III. OUR PROPOSED METHOD

Fig. 2 demonstrates the proposed HDR-GAN framework. Given three LDR images of different exposures (i.e., L_1 , L_2 and L_3 , sorted by their exposure biases) as input and L_2 as the reference image, our goal is to construct an HDR image. As in previous works [11]–[13], the input LDR images L_i ($i = \{1, 2, 3\}$) are first converted to the HDR domain (i.e., the value of each pixel is mapped to the range of $[0, 1]$), and then fused to produce an HDR image via the proposed deep model. Finally, following [11], the loss is measured by passing the generated HDR image through a differentiable tonemapper (i.e., μ -law), which is commonly used for range compression in audio processing:

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

where H represents an HDR image and μ controls the extent of compression.

The major challenge of fusing LDR images with different exposures to an HDR image is the presence of large motions in dynamic scenes, leading to the missing content and the misalignment among the LDR images. To encourage our model to hallucinate the missing details and thus to generate a high-quality HDR image, we propose to employ the GAN framework. In the past few years, GANs have been extensively studied and successfully applied to many image generation tasks. Similar to the standard GAN framework [32], [33], our model consists of a generator G and a discriminator D competing against each other in a two-player min-max game. In our problem, the generator tends to merge the LDR images and produce a high-quality HDR image. The discriminator, on the other hand, aims to classify the generated image and the ground-truth image as real or fake.

In the following subsections, we introduce the technical details of our proposed generator and discriminator.

A. The Generator

The purpose of the generator is to produce a high-quality HDR image \hat{H} , given the input LDR images L_i ($i = \{1, 2, 3\}$). As shown in Fig. 2, we follow existing practice [11], and feed the multi-exposed LDR images concatenated with their corresponding HDR versions as input \mathbf{X}_i ($i = \{1, 2, 3\}$) to the generator:

$$\mathbf{X}_i = \{L_i \oplus F_{\gamma}(L_i)\} = \{L_i \oplus L_i^\gamma/t_i\}, \quad (2)$$

where $F_{\gamma}(L_i)$ denotes the HDR version of L_i , γ denotes the Gamma correction parameter (with $\gamma = 2.2$) and t_i denotes the exposure time of L_i . \oplus denotes the concatenation operator. Hence, our generator network G can be briefly formulated as $\hat{H} = G(\mathbf{X}_1, \mathbf{X}_2, \mathbf{X}_3)$.

Due to the large motions of dynamic objects in the scene or the camera, the misalignment of the LDR images tend to produce misalignment of the extracted features, causing the model to generate undesired artifacts in the output HDR images. To address this concern, as shown in Fig. 2, we propose a novel generator network that is composed of specifically-designed modules, including multi-scale LDR encoding blocks and reference-aligned feature merging blocks, together with deep HDR supervision to enable implicitly alignment of the LDR images in the feature domain to reduce the artifacts of the resulting HDR images.

Multi-scale LDR encoder. For each input \mathbf{X}_i , we apply an individual encoder to extract its features in multiple scales, such that patch-level features of the LDR images can be learned for later fusion. Although prior works [12], [13] also introduce individual encoders, they do not fully exploit the multi-scale context information obtained from the encoders. Both of them adopt shallow encoders with two convolutional layers only, without considering multi-scale. In contrast, our encoders have three scales to process each input \mathbf{X}_i , and each scale is processed by one encoding block, i.e., a residual structure with two convolutional layers. The ablation study in Section IV-D demonstrates that our multi-scale LDR encoder performs the best.

Specifically, each LDR encoder of our network contains three consecutive downsampling encoding blocks (represented

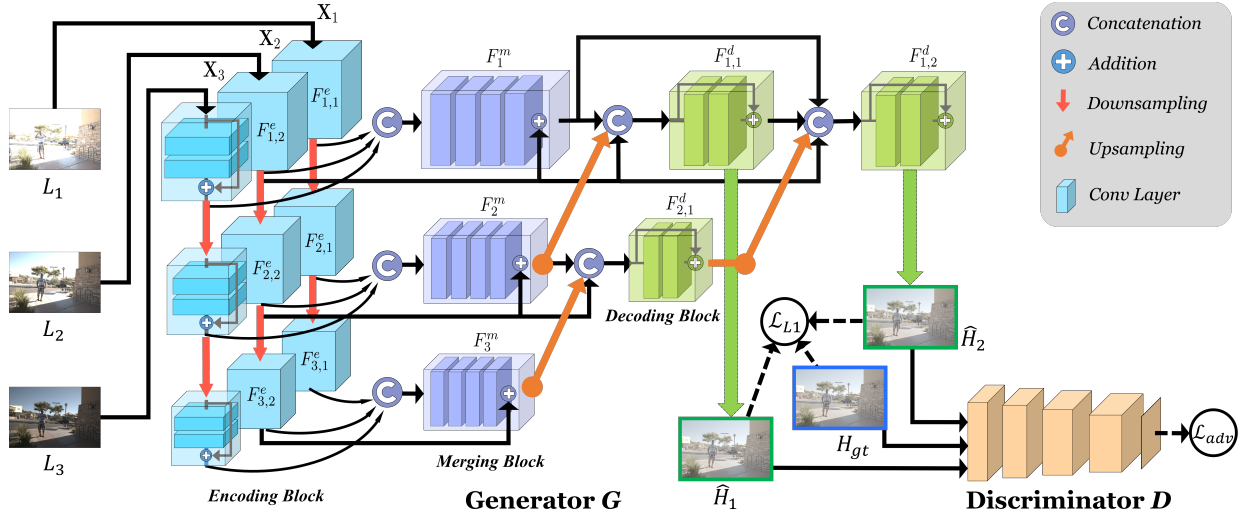


Fig. 2. **Illustration of the proposed framework.** Our framework is composed of two main networks: the generator G and the discriminator D . As illustrated, the three input LDR images L_i ($i \in \{1, 2, 3\}$) are fed into the generator network G to produce the output HDR image. G consists of the encoding blocks (i.e., the blue cubes) that extract features from the LDR images, the merging blocks (i.e., the light purple cubes) that align the LDR features, and the decoding blocks (i.e., the green cubes) that restore the aggregated features to larger scales. It produces two HDR images \hat{H}_1 and \hat{H}_2 for supervision with L1 loss, while \hat{H}_2 represents the output HDR image. On the other hand, the discriminator network D is employed to discriminate between the generated images and the ground-truth via an adversarial loss.

by the blue cubes from the top to bottom of Fig. 2), which help extract visual features in 1^{st} , 2^{nd} , and 3^{rd} scales from one input \mathbf{X}_i . Formally, we denote the encoding blocks of different scales as $F_{j,i}^e$, $j, i = \{1, 2, 3\}$, where i refers to the index of the input \mathbf{X} and j refers to the index of the network scale. Each encoding block is a two-layer residual structure without the pooling layers, followed by a convolutional layer with a 3×3 kernel size and stride 2 for downsampling the feature maps (shown as red arrows in Fig. 2). Hence, the encoded features at the j -th scale for \mathbf{X}_i , denoted as $\mathbf{E}_{j,i}$, is computed by sequentially passing \mathbf{X}_i through the encoding blocks from the top scale (i.e., $j = 1$) to the bottom scale (i.e., $j = 3$) as:

$$\begin{aligned} \mathbf{E}_{1,i} &= F_{1,i}^e(\mathbf{X}_i), \\ \mathbf{E}_{2,i} &= F_{2,i}^e(F_{1,i}^{down}(\mathbf{X}_i)), \\ \mathbf{E}_{3,i} &= F_{3,i}^e(F_{2,i}^{down}(F_{1,i}^{down}(\mathbf{X}_i))), \end{aligned} \quad (3)$$

where $F^{down}(\cdot)$ is implemented as a convolutional layer with stride 2.

Reference-aligned feature fusion. After extracting the features from the input LDR images, we fuse the features of different scales using the proposed merging blocks as shown in Fig. 2, which also help align the image content in the feature domain with a residual structure. In particular, the encoded features from three encoding blocks of the same scale are concatenated and transformed by a merging block, i.e., F_j^m ($j = \{1, 2, 3\}$), as:

$$\mathbf{M}_j = F_j^m(\{\mathbf{E}_{j,1} \oplus \mathbf{E}_{j,2} \oplus \mathbf{E}_{j,3}\}, \mathbf{E}_{j,2}), \quad (4)$$

where \mathbf{M}_j is the merged features of the j^{th} scale. \oplus denotes the concatenation operator. Hence, $\{\mathbf{E}_{j,1} \oplus \mathbf{E}_{j,2} \oplus \mathbf{E}_{j,3}\}$ represents the concatenation of the encoded features from multiple exposures in the j^{th} scale. In each merging block, we introduce a reference-based residual structure. According to Eq. 4, there is an additional input $\mathbf{E}_{j,2}$ (i.e., the encoded

features from the reference image) to the merging block F_j^m . $\mathbf{E}_{j,2}$ is directly fed to the end of the merging block to form a residual structure. The purpose of this design is to enforce the alignment of the encoded features by utilizing the features from the three LDR images to compensate the features from the reference image (i.e., $\mathbf{E}_{j,2}$). In essence, the features from the low and high exposures will be adapted to those of the median exposure (i.e., the reference image) to mitigate the misalignment problem. Represented by light purple cubes in Fig. 2, each merging block is composed of four dilated convolutional layers with 3×3 kernel size and a dilated factor of 2 to transform the concatenated features and allow $\mathbf{E}_{j,2}$ to skip-connect to the end of the block.

Deep HDR supervision. To further eliminate the artifacts in order to obtain high-fidelity results, we propose a deep HDR supervision scheme. Compared with prior deep supervision structures [16]–[18], [34], we upsample the intermediate merged features and concatenate them with those in the upper scales to produce the HDR images at the original spatial dimension using a series of decoding blocks. With the aggregated cross-scale features and the skip-connected encoded features from the reference image involved in the decoders, as demonstrated in the ablation study in Section IV-D, the deep supervision can further enforce the alignment of the merged features and thus our network is able to produce better quality HDR images.

Specifically, at the j^{th} scale, we concatenate the features from the reference image (i.e., $\mathbf{E}_{j,2}$), the merged features \mathbf{M}_j computed by Eq. 4, and the upsampled features of the lower scale $F^{up}(\mathbf{M}_{j+1})$ into feature maps with rich information. Note that $\mathbf{E}_{j,2}$, \mathbf{M}_j and $F^{up}(\mathbf{M}_{j+1})$ have all been aligned with the features from the reference image.

Recall that in the traditional encoder-decoder architecture (e.g., [34]), the encoded features of the deepest scale will be fed to the decoding blocks and upsampled sequentially. To

enhance the image generation quality, instead of supervising only the last decoding block as the standard decoders do, we also want to control the fused intermediate features of the network. To do this, besides the standard upsampling path (i.e., $F_3^m \rightarrow F_{2,1}^d \rightarrow F_{1,2}^d$ in Fig. 2), we introduce an auxiliary upsampling path (i.e., $F_2^m \rightarrow F_{1,1}^d$ in Fig. 2). While the standard upsampling path outputs an image \hat{H}_2 at the same dimension as the inputs, the auxiliary upsampling path aggregates the intermediate features and produces \hat{H}_1 also at the same size as the inputs.

Note that since the auxiliary upsampling path consists of fewer layers, it leads to faster convergence during training and its output is concatenated to the last decoding block, $F_{1,2}^d$. This upsampling path not only enhances the quality of the output images, but also benefits network learning, i.e., \hat{H}_2 can be complimentary to the incomplete details of \hat{H}_1 . In addition, the merged features at the top scale, F_1^m , is also consecutively passed through $F_{1,1}^d$ and $F_{1,2}^d$ for a thorough spatial transformation to further refine the features. Hence, the computation of each block of the decoder can be expressed as:

$$\begin{aligned} \mathbf{C}_{1,1} &= F_{1,1}^d(\mathbf{M}_1 \oplus \mathbf{E}_{1,2} \oplus F^{up}(\mathbf{M}_2)), \\ \mathbf{C}_{2,1} &= F_{2,1}^d(\mathbf{M}_2 \oplus \mathbf{E}_{2,2} \oplus F^{up}(\mathbf{M}_3)), \\ \mathbf{C}_{1,2} &= F_{1,2}^d(\mathbf{M}_1 \oplus \mathbf{E}_{1,2} \oplus \mathbf{C}_{1,1} \oplus F^{up}(\mathbf{C}_{2,1})), \end{aligned} \quad (5)$$

where $\mathbf{C}_{j,k}$ refers to the output of the decoding block $F_{j,k}^d$ ($k = \{1, \dots, 3-j\}$) and F^{up} denotes the upsampling layer. For each scale, we concatenate the features computed from the preceding layers with the upsampled features from the smaller scale to strengthen the feature representation. Besides, we also introduce the features from \mathbf{X}_2 (i.e., $\mathbf{E}_{j,2}$) that contains the reference image with median exposure to join the decoding blocks via skip connections. Furthermore, as shown in the green cubes of Fig. 2, each decoding block has an identical residual structure as the encoding blocks, which is followed by a nearest-neighbor interpolation layer.

For training, our model generates two HDR images \hat{H}_1 and \hat{H}_2 at the two output branches $\mathbf{C}_{1,1}$ and $\mathbf{C}_{1,2}$. Both are supervised by the ground-truth HDR image H_{gt} via an L_1 loss as:

$$\mathcal{L}_{L1} = \min_G (||\mathcal{T}(\hat{H}_1) - \mathcal{T}(H_{gt})||_1 + ||\mathcal{T}(\hat{H}_2) - \mathcal{T}(H_{gt})||_1), \quad (6)$$

where $\mathcal{T}(\cdot)$ is the differentiable tonemapper defined in Eq. 1. Thus, our deep HDR supervision scheme is able to eliminate the artifacts and improve the quality of the reconstructed HDR images.

B. The Discriminator

We adopt PatchGAN [33], which contains five convolutional layers, as the discriminator D . Unlike the standard discriminator, two generated HDR images, \hat{H}_1 and \hat{H}_2 , as well as the ground-truth H_{gt} are altogether passed into our discriminator. By discriminating \hat{H}_1 and \hat{H}_2 from H_{gt} , we can further encourage the generator to produce better quality HDR images, which will be demonstrated in Section IV-D. In addition, we also adopt the adversarial loss of [19] for stabilizing the training process and avoiding modal collapse.

Specifically, the output of the discriminator can be reshaped as an n -dimensional vector, i.e., $\mathbf{q} = D(\cdot) \in \mathbb{R}^n$, and then projected to a point \mathbf{p} on a hypersphere \mathbb{S}^n by the inverse of the stereographic projection, as:

$$\mathbf{p} = \left(\frac{2\mathbf{q}}{||\mathbf{q}||^2 + 1}, \frac{||\mathbf{q}||^2 - 1}{||\mathbf{q}||^2 + 1} \right). \quad (7)$$

Thus, instead of measuring the Euclidean distance between two points \mathbf{q} and \mathbf{q}' , we can measure the distance of any two projected points on the hypersphere by $d_s(\mathbf{p}, \mathbf{p}')$ as:

$$d_s(\mathbf{p}, \mathbf{p}') = \arccos \frac{||\mathbf{q}||^2 ||\mathbf{q}'||^2 - ||\mathbf{q}||^2 - ||\mathbf{q}'||^2 + 4\mathbf{q}\mathbf{q}' + 1}{(||\mathbf{q}||^2 + 1)(||\mathbf{q}'||^2 + 1)}. \quad (8)$$

In this GAN-based formulation, we can optimize the discriminator based on the distances of the mapping features from the generated images and those from the ground-truth images, with respect to the reference point $\mathbf{N} = [0, \dots, 0, 1]^T \in \mathbb{R}^n$, i.e., the north pole of the hypersphere. The formulation can be expressed as:

$$\begin{aligned} \mathcal{L}_{adv} &= \min_G \max_D \sum_r \mathbb{E}_{\mathbf{z}} [d_s^r(\mathbf{N}, D(\mathbf{z}))] \\ &\quad - \sum_r \mathbb{E}_{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3} [d_s^r(\mathbf{N}, D(G(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)))], \end{aligned} \quad (9)$$

where $G(\cdot)$ and $D(\cdot)$ refer to the generator and discriminator, respectively. $d_s^r(\cdot, \cdot)$ measures the r^{th} moment distance (we set $r = \{1, 2, 3\}$) on the hypersphere between $D(\cdot)$ and the north pole of the hypersphere (i.e., \mathbf{N}). In addition, \mathbf{z} is sampled from the distribution of the ground-truth images, while \mathbf{x}_i is sampled from the LDR images. To sum up, we train our HDR-GAN with the following hybrid loss:

$$\mathcal{L} = \mathcal{L}_{L1} + \lambda \cdot \mathcal{L}_{adv}, \quad (10)$$

where λ is a predefined constant that balances the loss terms, and we set it to 1 in our implementation.

IV. EXPERIMENTAL RESULTS

A. Datasets and Metrics

To evaluate the proposed method, we use the HDR dataset [11] for training and evaluation. It contains 74 image sets for training and 15 image sets for testing. For each training image set, three different LDR images are captured with exposure biases of $\{-2, 0, +2\}$ or $\{-3, 0, +3\}$. We also test our model on the datasets without ground-truth, including Sen's and Tursun's datasets [15], [35].

We use PSNR and SSIM as the metrics for evaluation. Specifically, we compute PSNR_μ and SSIM_μ between the fused HDR image (i.e., \hat{H}_2) and the ground truth image, both in the tonemapping domain by μ -law (Eq. 1) with $\mu = 5000$, which is consistent with the setting of the prior works [11]–[13]. We also compute PSNR_L and SSIM_L for comparison in the linear domain (i.e., HDR domain). Further, we compute the HDR-VDP-2 [36], which assesses the visibility and quality of the HDR images in different luminance conditions. Following [12], for the two parameters used to compute the HDR-VDP-2 scores, we set the diagonal display size to 24 inches, and the viewing distance to 0.5 meter.

TABLE I

COMPARISON WITH THE STATE-OF-THE-ART METHODS. THE BEST PERFORMANCE VALUES ARE INDICATED IN BOLD. [†] INDICATES THAT THE RESULTS ARE FROM [12] AND [14]. (ALL THESE RESULTS ARE THE HIGHER THE BETTER.)

Method	PSNR _μ	PSNR _L	SSIM _μ	SSIM _L	HDR-VDP-2
Sen <i>et al.</i> [15] [†]	40.800	38.110	0.9808	0.9721	59.38
Hu <i>et al.</i> [38] [†]	35.790	30.760	0.9717	0.9503	57.05
Kalantari <i>et al.</i> [11]	42.670	41.232	0.9888	0.9846	65.05
DeepHDR [12] [†]	41.650	40.880	0.9860	0.9858	64.90
AHDRNet [13]	43.631	41.143	0.9900	0.9702	64.61
NHDRNet [14] [†]	42.414	-	0.9887	-	61.21
Ours	43.922	41.572	0.9905	0.9865	65.45

B. Implementation and Details

We have implemented our model using Tensorflow and evaluated it on an NVIDIA Tesla V100 GPU with 16G memory. We adopt spectral normalization [37] instead of batch normalization in our model. All training images are at 1500×1000 resolution. Specifically, we sample patches of size 512×512 from the training set for training. To augment the training set, we randomly crop the training images and apply random rotation/flipping on the patches. The learning rate of our network is 10^{-4} initially, decays to 10^{-5} after 114,300 iterations, and then drops to 10^{-6} after 706,400 iterations. We train our model with a batch size of 2 for 1,215,000 iterations in total using the Adam solver, and the entire training process takes around a week. During inference, our model takes on average 0.29s and 8,833 MB memory to process a set of LDR images at 1440×960 resolution.

C. Comparison with the State-of-the-Arts

To evaluate our model, we compare it with the state-of-the-art HDR methods on the test images in the HDR dataset [11]. The state-of-the-art methods used for comparison include two patch-based methods, Sen *et al.* [15] and Hu *et al.* [38], and four deep-learning-based methods, Kalantari *et al.* [11], DeepHDR [12], AHDRNet [13], and NHDRNet [14]. Note that Kalantari *et al.* [11] apply optical flow to align the input images in the preprocessing stage. DeepHDR [12] requires the background of the input images to be aligned by homography transformation, while AHDRNet [13] and our method do not need to align the input images in the preprocessing stage.

Table I shows the results of the experiment. Note that the quantitative results for Sen *et al.* [15], Hu *et al.* [38], and DeepHDR [12] are from [12]; those for NHDRNet [14] are from [14]. On the other hand, we re-train the models from [11] and [13] and report the collected results. This is because we find that their reported results are slightly lower than the results that we obtain in running their models. Nevertheless, we can see that our method outperforms all the comparison methods on all metrics. Our top-performing results in the terms of PSNR_L and PSNR_μ implies the generation capability of our proposed method in the HDR domain and the tonemapping domain. Compared with the latest works [13], [14] that introduce attention schemes to relieve the LDR misalignment problem, their results in the terms of SSIM_μ and SSIM_L are not on par with those of our approach, indicating

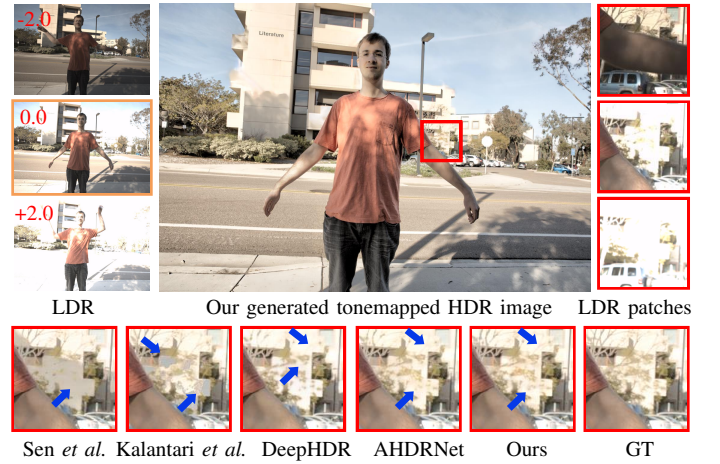


Fig. 3. An example from the dataset [11]. We compare a set of patches cropped from the tonemapped HDR images generated by the state-of-the-art methods. The blue arrows highlight the differences among the results of the comparison methods.

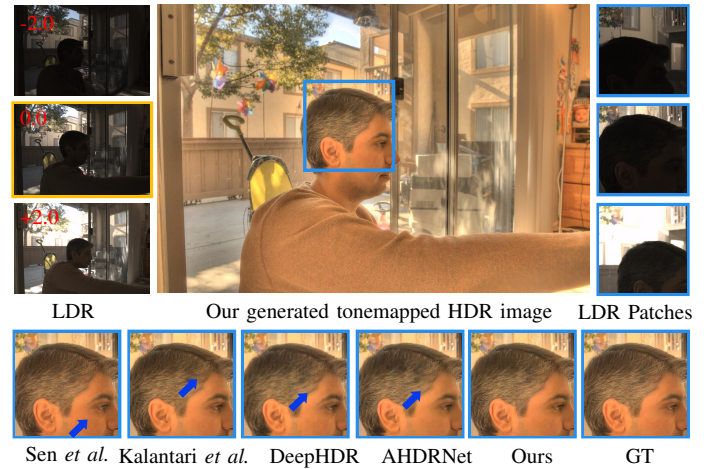


Fig. 4. Another example from the dataset [11]. We compare a set of patches cropped from the tonemapped HDR images generated by the state-of-the-art methods. The blue arrows highlight the differences among the results of the comparison methods.

that their models cannot consistently preserve the structural information of produced HDR images as ours. In addition, our results also demonstrate a healthy margin over other baselines on HDR-VDP-2 for the produced HDR image quality.

We also show several example scenarios with dynamic objects or camera motions from public datasets for visual comparison, as shown in Figs. 3, 4, 5 and 6. Note that the HDR images are tonemapped using *Photomatrix*¹, which is a tonemapper different from the one used in training. Figs. 3 and 4 show that our model is capable of recovering the image details, such as hair and sticks, from the LDR images, thanks to the LDR alignment in the feature domain. Specifically, Sen *et al.* [15] tends to introduce noise to the generated images (e.g., the face in Fig. 4), and the other methods may produce various degrees of blurry artifacts. On the contrary, our GAN-based method is able to generate more faithful information. Fig. 5 shows an image with a walking person and global

¹<https://www.hdrsoft.com>

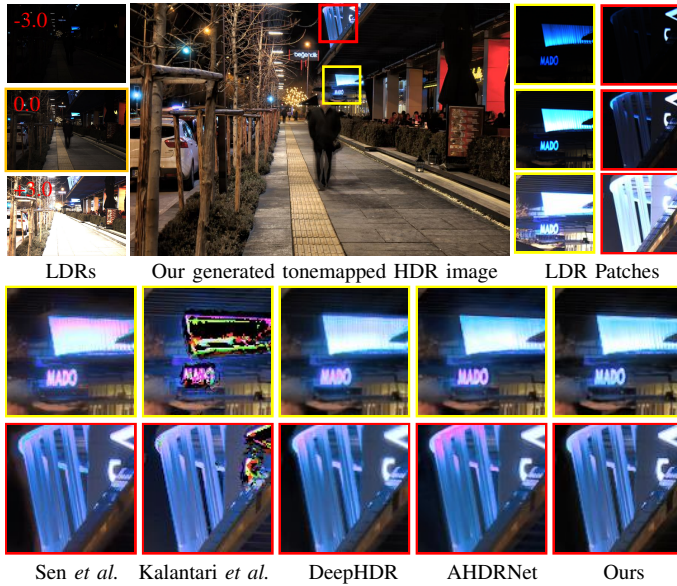


Fig. 5. An example from another dataset [35] (which does not provide ground-truth images). We qualitatively compare two sets of patches cropped from the tonemapped HDR images generated by the state-of-the-art methods.




Fig. 6. An example from yet another dataset [15] (which does not provide ground-truth images). We qualitatively compare one set of patches cropped from the tonemapped HDR images generated by the state-of-the-art methods. The red arrows highlight the differences among the results of the comparison methods.

camera motion, under a low-light condition. While the other methods may produce HDR images with distorted colors, which is not consistent with the reference image, due to the misalignment, our method still can well align LDR images and performs well in such a challenging scene. Fig. 6 shows a fast moving person, artifacts can be easily observed in the generated HDR images from all existing methods. In contrast, our result is much better, without the artifacts, indicating the effectiveness of our LDR feature alignment and the GAN-based paradigm of our model.

D. Ablation Study

In this subsection, we analyze the effectiveness of different parts of our model.

Network scale. To dissect the structure of our proposed generator, we first analyze the scale of its network structure. We introduce several variants for comparison, as shown in Fig. 7, including a two-scale version (Fig. 7(a)), a four-scale

TABLE II
ABLATION STUDY OF DIFFERENT VARIANTS OF THE PROPOSED NETWORK STRUCTURE ON THE FOUR METRICS. D.S. INDICATES WHETHER THE DEEP SUPERVISION IS ADOPTED IN A MULTI-SCALE NETWORK. THE VALUES HIGHLIGHTED IN RED AND BLUE INDICATE THE BEST AND THE SECOND-BEST PERFORMANCES, RESPECTIVELY.

Structure	D.S.	Output	PSNR _μ	PSNR _L	SSIM _μ	SSIM _L
Two-scale	-	\hat{H}_1	43.710	41.139	0.9904	0.9855
Four-scale	✓	\hat{H}_1	43.847	41.318	0.9906	0.9860
		\hat{H}_2	43.901	41.416	0.9906	0.9865
		\hat{H}_3	43.833	41.363	0.9906	0.9867
w/o Merging blocks	-	\hat{H}_1	43.595	41.009	0.9900	0.9854
Ours	✓	\hat{H}_1	43.787	41.223	0.9904	0.9853
		\hat{H}_2	43.922	41.572	0.9905	0.9865
	×	\hat{H}_2	43.838	41.381	0.9905	0.9857

version (Fig. 7(b)), and our proposed three-scale network (Fig. 7(d)). Note that the two-scale network contains only one output branch for supervision (i.e., \hat{H}_1), while the four-scale network contains three output branches (i.e., \hat{H}_1 , \hat{H}_2 and \hat{H}_3). Table II shows the corresponding results (i.e., first, second and fourth groups of results). Intuitively, a shallower network may not be able to capture sufficient multi-scale features for the task, while a deeper network can cause redundancy, network learning difficulties or even overfitting. We can observe from Table II that the two-scale network does not perform well, while the deeper network has slightly lower performances, compared with our proposed three-scale network. It is worth noting that the best performing output branch of the four-scale network is the second one, which implies that the last output branch of the network may be overfitted during training. From these results, we may conclude that the three-scale architecture is the optimal one.

Fig. 8 shows two patches of an example image, and the results of these two patches from the three variants. In this challenging scene, the background is over-exposed and partially occluded due to the motions of the two persons. We can see that both the two-scale network (Fig. 8(b)) and the four-scale network (Fig. 8(c)) have obvious artifacts. In contrast, our model (Fig. 8(d)) produces much better results.

Merging blocks. We then study the effectiveness of the merging blocks. As mentioned, to align the features from multiple exposures and preserve the details of the generated images, we densely align the features via the residual merging blocks. To validate this, we compare our network with the network without the merging blocks as shown in Fig. 7(c). The comparison results are also shown in Table II. The model w/o merging blocks (third groups of results in Table II) has a performance significantly lower than the proposed model (fourth groups of results in Table II), in terms of PSNRs. We can also observe from the visual results in Fig. 8(a) that without the merging blocks, the model fails to align the three input images and is thus unable to recover the contents well.

Deep HDR supervision. To demonstrate the effectiveness of our deep HDR supervision, we further conduct several experiments for a thorough analysis. First, we evaluate the generated HDR results with and without deep supervision on

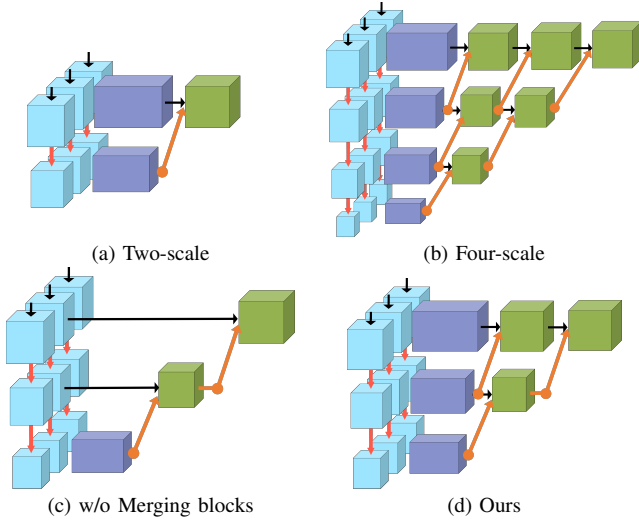


Fig. 7. Different variants of our network structure, including (a) two-scale and (b) four-scale, (c) the network structure without the merging blocks, and (d) our proposed model.

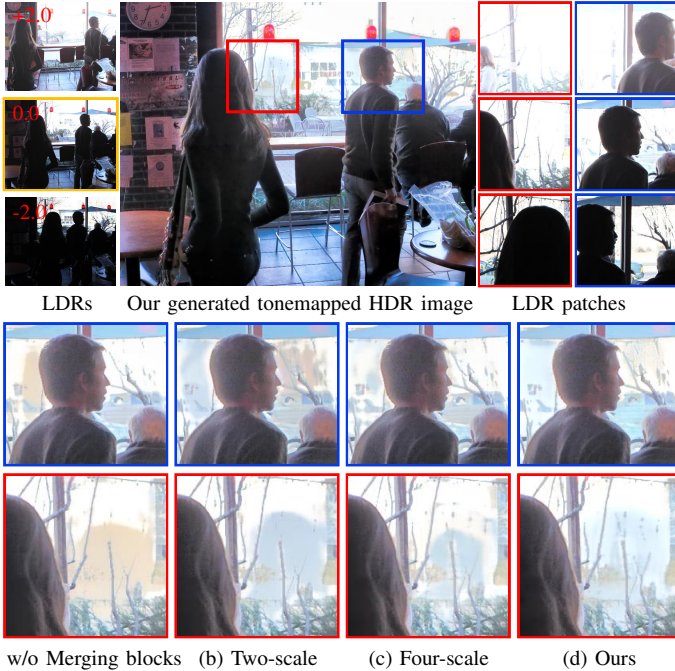


Fig. 8. An example to compare the visual results of different variants of our network structure.

the metrics. In addition, we also compare these results with the results from the network trained using only the second output branch (i.e., \hat{H}_1 is discarded and only \hat{H}_2 is used in the loss function). The last groups of results in Table II shows these results. As observed, in our proposed model, the result of \hat{H}_2 outperforms that of \hat{H}_1 , since the information from \hat{H}_1 helps further refine \hat{H}_2 by compensating with more details to the generated \hat{H}_2 . For reference, given the model training with \hat{H}_2 only, it performs worse than the one with multiple supervisions. We show a visual example in Fig. 9, which demonstrates the results from different output branches. We can see the best result is produced when using both \hat{H}_1 and \hat{H}_2 in the loss function and with \hat{H}_2 as the output, which is



Fig. 9. We show example HDR images generated by using (a) using \hat{H}_2 as the model output and without using \hat{H}_1 in model training, (b) using \hat{H}_1 as the model output, and (c) using \hat{H}_2 as the model output.




Fig. 10. Training losses of the two output branches (\hat{H}_1 and \hat{H}_2).

consistent with our quantitative results.

Second, we show the training losses of our proposed model in Fig. 10. The gray curve indicates the training loss of \hat{H}_1 , while the red curve indicates the training loss of \hat{H}_2 . We can see from the results that when the number of iterations is small ($< 40,000$), \hat{H}_1 performs better. This is because \hat{H}_1 has a shallower structure. However, with a higher number of iterations ($> 70,000$), \hat{H}_2 performs comparatively better.

Discriminator. To evaluate the effectiveness of our discriminator D for generating HDR images, we first compare the results from our model without the discriminator D . In particular, as observed in Table III, comparing the models with and without D , the performance of our model with D is apparently boosted. It is interesting to mention that our generator network itself (i.e., without introducing the discriminator nor the GAN loss) can achieve comparable performance with the state-of-the-art methods. Specifically, we can observe from Tables I

TABLE III

ABLATION STUDY ON THE EFFECTIVENESS OF OUR DISCRIMINATOR D . WE FIRST COMPARE AGAINST THE FRAMEWORK WITHOUT THE DISCRIMINATOR (I.E., ‘w/o D ’). WE ALSO COMPARE WITH THE DISCRIMINATOR WHICH TAKES ONLY ONE GENERATED HDR IMAGE (\hat{H}_1 OR \hat{H}_2) AND THE GROUND-TRUTH IMAGE AS INPUTS. THE VALUES HIGHLIGHTED IN RED AND BLUE INDICATE THE BEST AND THE SECOND-BEST PERFORMANCE, RESPECTIVELY.

Structure	Output	PSNR $_{\mu}$	PSNR $_L$	SSIM $_{\mu}$	SSIM $_L$
w/o D	\hat{H}_1	43.699	41.005	0.9903	0.9857
	\hat{H}_2	43.729	41.013	0.9903	0.9858
w/ D (w/o \hat{H}_2 in \mathcal{L}_{adv})	\hat{H}_1	43.634	41.132	0.9902	0.9854
	\hat{H}_2	43.693	41.085	0.9903	0.9857
w/ D (w/o \hat{H}_1 in \mathcal{L}_{adv})	\hat{H}_1	43.535	40.991	0.9901	0.9851
	\hat{H}_2	43.756	41.281	0.9903	0.9861
Ours	\hat{H}_1	43.787	41.223	0.9904	0.9853
	\hat{H}_2	43.922	41.572	0.9905	0.9865

and III that our trained generator without GAN obtains a better result in terms of PSNR $_{\mu}$ (i.e., 43.729) than AHDRNet (i.e., 43.631), while our model with GAN obtains an even better result (i.e., 43.922).

In addition, we also analyze the effectiveness of using multiple generated HDR images in the discriminator. Recall that during training, our discriminator D receives two deep supervised results from the generator (i.e., \hat{H}_1 and \hat{H}_2) as well as the ground-truth image (i.e., H_{gt}) as inputs. These two deep supervised results are supposed to be similar, but not exactly the same. To evaluate if it is necessary to have both of them as inputs to the discriminator, we train our GAN-based model with \hat{H}_1 or \hat{H}_2 only. In other words, according to Eq. 9, only the first output \hat{H}_1 or the second output \hat{H}_2 from $G(\cdot)$ is used in the adversarial loss \mathcal{L}_{adv} . Table III shows the results, denoted as “w/ D (w/o \hat{H}_2 in \mathcal{L}_{adv})” and “w/ D (w/o \hat{H}_1 in \mathcal{L}_{adv})”. From these results, we can observe that sending two outputs to the discriminator helps improve the network generation ability slightly. This is because D needs to discriminate not only \hat{H}_2 but also \hat{H}_1 from the ground-truth, which thus encourages the model to improve the quality of the generated HDR images.

V. CONCLUSION AND FUTURE WORKS

In this paper, we propose a novel GAN-based HDR model, *HDR-GAN*, to address the problems caused by large object motions in the scene. Given three multi-exposed LDR images as input, our model generates an HDR image. By incorporating the adversarial learning scheme, our method is able to produce faithful information in the regions with missing contents. In addition, we also propose a novel generator network, with reference-based residual merging blocks for aligning large object motions in the feature domain, and a deep HDR supervision scheme for eliminating artifacts of the reconstructed HDR images. Our extensive experiments demonstrate that the proposed model can obtain state-of-the-art reconstruction performance, compared with the prior HDR methods.

Generally, HDR imaging is still a challenging research problem, especially when the reference image of the input LDR images suffer from significant quality degeneration, e.g.,

severe missing content due to under- or over-saturation. Thus, existing datasets may not be sufficient to push forward the progress of HDR techniques, since the HDR dataset [11] only contains images from a limited number of less challenging scenes and its performance measures for the state-of-the-art methods are hard to be further improved, while the other datasets do not provide ground-truth. As a future work, it will be necessary to establish a larger HDR dataset with more diverse and challenging scenes. Besides, since the resolution of photos taken by the latest electronic products become very high, directly reconstructing HDR images from such high-resolution images is a formidable task for the off-the-shelf CNN-based models due to the limited memory resources. In the future, it is needed to investigate a model that can efficiently and effectively perform HDR imaging for high-resolution images.

REFERENCES

- [1] S. K. Nayar and T. Mitsunaga, “High dynamic range imaging: Spatially varying pixel exposures,” in *CVPR*, 2000.
- [2] J. Tumblin, A. K. Agrawal, and R. Raskar, “Why I want a gradient camera,” in *CVPR*, 2005.
- [3] P. E. Debevec and J. Malik, “Recovering high dynamic range radiance maps from photographs,” in *Conference on Computer Graphics & Interactive Techniques*, 1997.
- [4] S. Mann and R. W. Picard, “On being ‘undigital’ with digital cameras: Extending dynamic range by combining differently exposed pictures,” in *IS&T’s Annual Conference*, 1995.
- [5] M. Granados, B. Ajdin, M. Wand, C. Theobalt, and H. P. A. Lensch, “Optimal HDR reconstruction with linear digital cameras,” in *CVPR*, 2010.
- [6] E. Reinhard, G. Ward, S. Pattanaik, and P. E. Debevec, *High dynamic range imaging: Acquisition, display, and image-based lighting*, 2005.
- [7] K. Myszkowski, R. Mantiuk, and G. Krawczyk, “High dynamic range video,” *ACM TOG*, 2008.
- [8] H. Zimmer, A. Bruhn, and J. Weickert, “Freehand HDR imaging of moving scenes with simultaneous resolution enhancement,” *Computer Graphics Forum*, 2011.
- [9] T. Grosch, “Fast and robust high dynamic range image generation with camera and object movement,” in *IEEE International Conference of Vision, Modeling and Visualization*, 2006.
- [10] K. Jacobs, C. Loscos, and G. Ward, “Automatic high-dynamic range image generation for dynamic scenes,” *IEEE Computer Graphics & Application*, 2008.
- [11] N. K. Kalantari and R. Ramamoorthi, “Deep high dynamic range imaging of dynamic scenes,” *ACM TOG*, 2017.
- [12] S. Wu, J. Xu, Y.-W. Tai, and C.-K. Tang, “Deep high dynamic range imaging with large foreground motions,” in *ECCV*, 2018.
- [13] Q. Yan, D. Gong, Q. Shi, A. van den Hengel, C. Shen, I. Reid, and Y. Zhang, “Attention-guided network for ghost-free high dynamic range imaging,” in *CVPR*, 2019.
- [14] Q. Yan, L. Zhang, Y. Liu, Y. Zhu, J. Sun, Q. Shi, and Y. Zhang, “Deep HDR imaging via a non-local network,” *IEEE TIP*, 2020.
- [15] P. Sen, N. K. Kalantari, M. Yaesoubi, S. Darabi, D. B. Goldman, and E. Shechtman, “Robust patch-based HDR reconstruction of dynamic scenes,” *ACM TOG*, 2012.
- [16] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *CVPR*, 2015.
- [17] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, “Unet++: A nested u-net architecture for medical image segmentation,” in *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support*, 2018.
- [18] K. Sun, B. Xiao, D. Liu, and J. Wang, “Deep high-resolution representation learning for human pose estimation,” in *CVPR*, 2019, pp. 5693–5703.
- [19] S. W. Park and J. Kwon, “Sphere generative adversarial network based on geometric moment matching,” in *CVPR*, 2019.
- [20] N. Barakat, A. N. Hone, and T. E. Darcie, “Minimal-bracketing sets for high-dynamic-range image capture,” *IEEE TIP*, 2008.


- [21] S. W. Hasinoff, F. Durand, and W. T. Freeman, "Noise-optimal capture for high dynamic range photography," in *CVPR*, 2010.
- [22] K. Seshadrinathan, S. H. Park, and O. Nestares, "Noise and dynamic range optimal computational imaging," in *ICIP*, 2012.
- [23] R. Pourreza-Shahri and N. Kehtarnavaz, "Exposure bracketing via automatic exposure selection," in *ICIP*, 2015.
- [24] K. R. Prabhakar, V. S. Srikar, and R. V. Babu, "DeepFuse: A deep unsupervised approach for exposure fusion with extreme exposure image pairs," in *ICCV*, 2017, pp. 4724–4732.
- [25] P. v. Beek, "Improved image selection for stack-based HDR imaging," *Electronic Imaging*, 2019.
- [26] G. Eilertsen, J. Kronander, G. Denes, R. K. Mantiuk, and J. Unger, "HDR image reconstruction from a single exposure using deep CNNs," *ACM TOG*, 2017.
- [27] S. Lee, G. Hwan An, and S.-J. Kang, "Deep recursive HDRI: Inverse tone mapping using generative adversarial networks," in *ECCV*, 2018.
- [28] Z. Khan, M. Khanna, and S. Raman, "FHDR: HDR image reconstruction from a single LDR image using feedback network," in *Global SIP*, 2019.
- [29] Y.-L. Liu, W.-S. Lai, Y.-S. Chen, Y.-L. Kao, M.-H. Yang, Y.-Y. Chuang, and J.-B. Huang, "Single-image HDR reconstruction by learning to reverse the camera pipeline," in *CVPR*, 2020.
- [30] Q. Yan, D. Gong, P. Zhang, Q. Shi, J. Sun, I. Reid, and Y. Zhang, "Multi-scale dense networks for deep high dynamic range imaging," in *WACV*, 2019.
- [31] S.-Y. Chen and Y.-Y. Chuang, "Deep exposure fusion with deghosting via homography estimation and attention learning," in *ICASSP*, 2020.
- [32] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *NeurIPS*, 2014.
- [33] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *CVPR*, 2017.
- [34] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and Computer-assisted Intervention*, 2015.
- [35] O. Tursun, A. Akyüz, A. Erdem, and E. Erdem, "An objective deghosting quality metric for HDR images," *Computer Graphics Forum*, 2016.
- [36] R. Mantiuk, K. J. Kim, A. G. Rempel, and W. Heidrich, "HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions," in *ACM TOG*, 2011.
- [37] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," in *ICLR*, 2018.
- [38] J. Hu, O. Gallo, K. Pulli, and X. Sun, "HDR deghosting: How to deal with saturation?" in *CVPR*, 2013.



Yuzhen Niu received her Ph.D. degree in Computer Science from Shandong University of China in 2010. She was a Post-Doctoral Researcher with the Department of Computer Science, Portland State University, Portland, OR. She is currently a Professor with the College of Mathematics and Computer Science, Fuzhou University, China. Her current research interests include image and video processing, computer vision, artificial intelligence, and multimedia.



Jianbin Wu is pursuing a master degree of Computer Software and Theory at Fuzhou University, Fuzhou, China since 2018. His research interests include computational image processing, deep learning, and computer vision.



Wenxi Liu is an Associate Professor in the College of Mathematics and Computer Science, Fuzhou University. He obtained his Ph.D. degree from City University of Hong Kong. His research interests include computer vision, robot vision, and image processing.



Wenzhong Guo received the B.S. and M.S. degrees in computer science and the Ph.D. degree in communication and information systems from Fuzhou University, Fuzhou, China, in 2000, 2003, and 2010, respectively. He is currently a Full Professor with the College of Mathematics and Computer Science, Fuzhou University. He leads the Network Computing and Intelligent Information Processing Laboratory, which is a key Lab of Fujian Province, China. His research interests include mobile computing and evolutionary computation.



Rynson W.H. Lau received his Ph.D. degree from University of Cambridge. He was on the faculty of Durham University and is now with City University of Hong Kong.

Rynson serves on the Editorial Board of Computer Graphics Forum, and Computer Animation and Virtual Worlds. He has served as the Guest Editor of a number of journal special issues, including ACM Trans. on Internet Technology, IEEE Trans. on Multimedia, IEEE Trans. on Visualization and Computer Graphics, and IEEE Computer Graphics & Applications. He has also served in the committee of a number of conferences, including Program Co-chair of ACM VRST 2004, ACM MTDL 2009, IEEE U-Media 2010, and Conference Co-chair of CASA 2005, ACM VRST 2005, ACM MDI 2009, ACM VRST 2014. Rynson's research interests include computer graphics and computer vision.