

Single-Image HDR Reconstruction by Multi-Exposure Generation

Phuoc-Hieu Le¹ Quynh Le^{1,2} Rang Nguyen¹ Binh-Son Hua¹

¹VinAI Research ²University of California San Diego




Figure 1: The result from our weakly supervised single-image HDR reconstruction. Deep Recursive HDRI [18] and DrTMO [6] produce artifacts in saturated regions. Our method outperforms both previous methods with a more visually pleasing HDR.

Abstract

High dynamic range (*HDR*) imaging is an indispensable technique in modern photography. Traditional methods focus on *HDR* reconstruction from multiple images, solving the core problems of image alignment, fusion, and tone mapping, yet having a perfect solution due to ghosting and other visual artifacts in the reconstruction. Recent attempts at single-image *HDR* reconstruction show a promising alternative: by learning to map pixel values to their irradiance using a neural network, one can bypass the align-and-merge pipeline completely yet still obtain a high-quality *HDR* image. In this work, we propose a weakly supervised learning method that inverts the physical image formation process for *HDR* reconstruction via learning to generate multiple exposures from a single image. Our neural network can invert the camera response to reconstruct pixel irradiance before synthesizing multiple exposures and hallucinating details in under- and over-exposed regions from a single input image. To train the network, we propose a representation loss, a reconstruction loss, and a perceptual loss applied on pairs of under- and over-exposure images and thus do not require *HDR* images for training. Our experiments show that our proposed model can effectively reconstruct *HDR* images. Our qualitative and quantitative results show that our method achieves state-of-the-art performance on the DrTMO dataset. Our code is available at https://github.com/VinAIResearch/single_image_hdr.

1. Introduction

Cameras are optical instruments designed to mimic the *human visual system* (HVS). They can capture the surrounding environment as close as what our eyes can observe in general conditions. Unfortunately, an essential but challenging factor for the camera to reproduce compared to the human visual system is the dynamic range. Remarkably, dynamic ranges captured by a camera and our eyes are not the same. A consumer-grade camera can only capture images with relatively *low dynamic ranges* (LDR images) while our eyes can perceive very *high dynamic ranges* (HDR) [2, 7]. The images captured by such cameras often result in over-exposed regions with many saturated details.

HDR imaging is a technique in modern photography to reproduce a higher dynamic range in a photograph so that more details of the bright and dark regions can be retained in the image. *HDR* images are a more faithful reproduction of a scene, being closer to HVS than traditional (low) dynamic range images. Beyond photography, *HDR* has found applications in image-based lighting, *HDR* display design, and computer vision downstream tasks. As the need for *HDR* imaging becomes prevalent, techniques for *HDR* imaging reconstruction are requisite.

Unfortunately, acquiring *HDR* images is a challenging task. To reconstruct an *HDR* image, one typically needs special hardware like a camera system with *HDR* support; else, one has to capture multiple LDR images and reconstruct *HDR* computationally. The latter approach is more

popular as its theories, and best practices have been well understood and widely implemented on consumer devices such as smartphones [10]. A common technique is to reconstruct HDR with multiple exposure images, where each exposure captures details in certain dynamic ranges. Notwithstanding such adoption, multi-exposure-based HDR suffers from visual artifacts due to object motion at capture. Computational photography methods on HDR reconstruction mainly deal with approaches that can mitigate these artifacts.

With deep learning, reconstructing an HDR image from a single LDR image becomes a plausible solution to resolve visual artifacts caused by motion. In principle, we could generate a sequence of well-aligned images with different exposures, which can be fused by conventional methods [4, 23] to generate an HDR image. Recent work [6, 17, 18, 19] has made significant progress in this direction. However, these approaches are designed in a supervised learning manner that requires input images with corresponding ground truth. These approaches do not explicitly handle the missing-detail issue in saturated regions (as shown in Fig. 1). In this work, we propose a novel weakly supervised learning method that utilizes only low-dynamic range images from the same exposure stack for training HDR reconstruction. The basic idea is to learn to generate multiple exposure images from a single image by inverting the camera response and hallucinating missing details using neural networks. Our main contributions are:

- A novel end-to-end trainable neural network that can generate an arbitrary number of different exposure images for HDR reconstruction. Our network is designed with weakly supervised learning that only uses multiple exposures for training, and thus can relax the requirement of obtaining ground truth HDR images;
- An objective function that utilizes pixel irradiance for supervising the network using only multiple exposure images without the need for ground truth HDR image;
- Comprehensive quantitative and qualitative evaluations with results showing that the proposed framework is comparable to existing models.

We will release our implementation, evaluation code, and pre-trained model upon publication.

2. Related Work

High-quality HDR from Multiple Exposures. HDR reconstruction is a long-studied problem in computer vision. The typical approach to this problem is to reconstruct HDR from multiple exposures as suggested by Debevec and Malik [4] or Mertens *et al.* [23]. The aforementioned methods, however, often fail to reconstruct the desired HDR image properly, leading to artifacts, ghosting, and tearing in the

Approach	Method	Dataset available	Training
Direct Input: LDR Output: HDR	Eilertsen <i>et al.</i> [5]	No	Supervised
	Yang <i>et al.</i> [43]	No	Supervised
	Moriwaki <i>et al.</i> [25]	No	Supervised
	Marnierides <i>et al.</i> [22]	No	Supervised
	Khan <i>et al.</i> [14]	No	Supervised
	Liu <i>et al.</i> [21]	Yes	Supervised
	Santos <i>et al.</i> [36]	No	Supervised
Indirect Input: LDR Output: LDRs	Endo <i>et al.</i> [6]	Yes	Exposure stack
	Lee <i>et al.</i> [17]	No	Exposure stack
	Lee <i>et al.</i> [18]	No	Exposure stack
	Lee <i>et al.</i> [19]	No	Exposure stack
	Ours	Yes	Exposure pair

Table 1: Summary of single-image HDR reconstruction.

final HDR image, especially when motion is introduced in the scene. For a long time, research in HDR reconstruction was focused on mitigating such artifacts [10, 12, 37, 38].

In the modern era of deep learning, Kalantari *et al.* [13] proposed the first learning-based method for HDR image reconstruction for dynamic scenes, which performs image alignment and merging with a convolutional neural network (CNN). Later work [29, 30, 31] followed the previous pipeline but replaced the conventional optical flow in the alignment step with CNN. Others opt for an end-to-end network [40, 41, 42] or a generative adversarial neural network (GAN) [20, 27] to solve this problem. While these methods, as mentioned earlier, can produce high-quality HDR images, eliminating ghosting artifacts is still a challenging problem in the multiple exposure pipeline.

Single-Image HDR Reconstruction. Using a single image for HDR reconstruction is, therefore, beneficial in that the misalignment problem can be circumvented. Eilertsen *et al.* [5] proposed to use CNN to predict the missing information lost in saturated regions caused by sensor saturation. Differing from the previous work, Yang *et al.* [43] enriched details in LDR images by using a CNN that first recovered an HDR image with missing details. Then it learned a tone mapping function that mapped from HDR to the LDR domain with the retrieved details. Other ideas, such as employing a hybrid loss [25], combining local and global features [22] or using Feedback Network [14], have been proposed in an attempt to output more realistic results. Recently, Liu *et al.* [21] decomposed the HDR imaging problem into three sub-tasks based on modeling the reversed, HDR-to-LDR image formation pipeline: dynamic range clipping, non-linear mapping, and quantization. Note that a similar idea [1, 44] to reverse the camera pipeline has been applied in the denoising task. Santos *et al.* [36] suggested a feature masking mechanism to guide the network to focus on the valid information in the well-exposed regions rather than saturated ones to avoid causing ambiguity during training the CNN. Furthermore, their work also suggested that pre-training with inpainting can help the network synthesize visually pleasing

textures in the saturated areas. The previous work typically extends the dynamic range of the input by *directly* predicting details in the overexposed regions caused by sensor saturation. These methods require *ground truth HDRs* available in the learning process to achieve that goal.

An *indirect* way to reconstruct HDR from a single image is via the prediction of multiple exposure images. The final HDR photo is then reconstructed from the inferred bracketed LDR images. The benefit of this approach is that it allows more fine-grained control of the details by having the low- and high-exposure generation in separate processes. This idea was first explored by Endo *et al.* [6], where they used two neural networks to infer up- and down-exposure images from an LDR image with medium exposure. Similarly, Lee *et al.* [17] later proposed a single model containing six sub-networks in a chaining structure to infer the bracketing images sequentially. With the input at a middle exposure value (EV_0), their model can infer $EV \pm 1, \pm 2, \pm 3$ effectively as the network goes deeper. As the number of synthesized bracketed images is fixed along with the exposure of each image, any attempts to overcome these limitations may require re-training the network. To overcome the mentioned issue, Lee *et al.* [18] defined two neural networks representing the relationship between images with relative EVs. The proposed structure can scale well with the number of generated images without the need to re-train or to add more sub-networks. Following that, Lee *et al.* [19] improved the predecessor work by using two conditional GAN structures [24] to generate a multi-exposure stack recursively. Although the mentioned frameworks can synthesize plausible multi-exposure stacks, it still has limitations since it has neither more granular control over the output exposure nor takes into account the image formation pipeline. Our method is based on this indirect approach in that our network predicts multiple exposure images granularly by inverting the physical image formation pipeline.

We summarize all methods for single-image HDR reconstruction in Tab. 1, where our method is based on weak supervision from multi-exposure images. Note that while the method of Ram *et al.* [32] is unsupervised, our method is different from that as we learn to generate multi-exposure stacks, whereas Ram *et al.*'s method is used to fuse them.

3. Our Approach

3.1. Problem formulation

In this section, we propose our method for HDR image reconstruction. The basic idea is to let the network learn to generate multiple exposures from a single input image, and the HDR can be reconstructed from the generated exposures following the conventional HDR pipeline. Let us begin with our camera pipeline for image formation (Fig. 2). We generally model an image I from the in-camera image

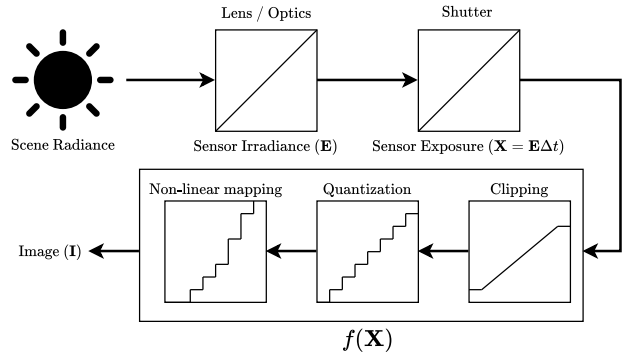


Figure 2: We model the traditional image formation pipeline by a non-linear function $f(X)$ where $X = E\Delta t$ is the sensor exposure formed by the product of sensor irradiance E and exposure time t . Our method learns to invert this function by using CNNs in an end-to-end fashion.

processing pipeline as a function $f(X)$ that transforms the scene irradiance E integrated over exposure time Δt . The light measurement $X = E\Delta t$ is linear and based on the physics model of light transport, and the function $f(X)$ is the composition of all non-linearities that occurred in the in-camera signal processing, such as camera response function, clipping, quantization, and other non-linear mappings. In this model, the sensor irradiance E captures the high dynamic range signals of the scene, and the image I represents the low dynamic range (LDR) signal for display. Note that we assume the noise in the process is insignificant. Thus we decide to exclude it from the pipeline.

To perform HDR reconstruction, our goal is to invert this image formation model to recover sensor irradiance E from image I so that $E = f^{-1}(I) / \Delta t$. This means we have to invert the non-linearities captured in $f(X)$. Unfortunately, this is an ill-posed and challenging problem because some steps in $f(X)$ are irreversible, e.g., clipping, while the camera response function varies per camera and is usually considered proprietary. To address this problem, we opt for data-driven approaches, and we propose to use CNNs to learn the pipeline inversion.

3.2. Proposed network

Our HDR supervision is based on pairs of images taken with different exposures. We assume that the images are taken with the same camera in the same scene, so the images share the same underlying scene irradiance, which is also the same assumption as Grossberg and Nayar's method [8]. Let the multiple exposures be $\{I_i\}$ with $I_i = f(X_i) = f(E\Delta t_i), \forall i = 1, \dots, n$ where Δt_i is the exposure time of a corresponding image I_i and sensor exposure $X_i = E\Delta t_i$ to supervise the neural network. Particularly, for every pair of low- and high-exposure (I_1, I_2) from the same sensor irradiance E and mapping function f with corresponding

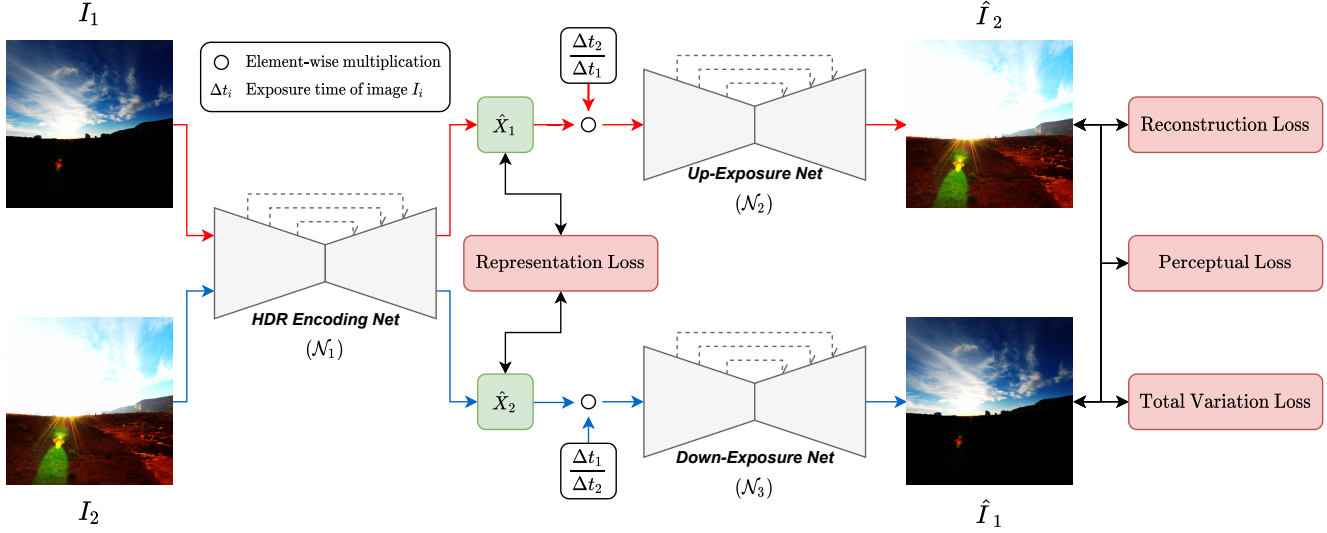


Figure 3: Training pipeline of our proposed framework. Given a pair of images in two different exposures, we predict latent invariant representation from the exposures by enforcing the exposure pair (\hat{X}_1, \hat{X}_2) to have the same representation when scaled by a factor (network \mathcal{N}_1). This representation can then be scaled and passed to Up/Down-Exposure Net (\mathcal{N}_2 and \mathcal{N}_3) to reconstruct different exposure images.

exposure time $(\Delta t_1, \Delta t_2)$ where $\Delta t_2 > \Delta t_1$ as input, we predict the up- and down-exposure \hat{I}_2 and \hat{I}_1 for I_1 and I_2 , respectively. This task guides the network to generate \hat{I}_1 and \hat{I}_2 such that they match well the input exposure I_1 and I_2 . Mathematically, the relation between I_1 and I_2 can be written as:

$$X_i = E\Delta t_i = f^{-1}(I_i) \quad (1)$$

where $i \in \{1, 2\}$. Given X_i , we can then scale it accordingly and generate a different exposure image I_j :

$$I_j = f(X_j) = f(E\Delta t_j) = f\left(X_i \frac{\Delta t_j}{\Delta t_i}\right) \quad (2)$$

where $j \in \{2, 1\}$. Figure 3 shows the overall structure of our proposed method.

Network overview. Our proposed network consists of two stages. The first stage is backward mapping, where we use our **HDR Encoding Net** (\mathcal{N}_1) to transform input image I_i into X_i , a suitable representation for the image’s sensor exposure at exposure time Δt_i in latent space. An appropriate factor can scale the representation X_i , as shown in Eq. (2) to get the sensor exposure at a shorter or longer exposure time $\Delta t_{i\pm 1}$.

The next stage is forward mapping, where we want to map the scene irradiance X_i into a pixel value after altering its exposure time Δt_i to generate a new image with a different exposure. In this stage, it requires to hallucinate details in saturated regions. Due to the different nature of under- and over-exposed images, we propose to use two sub-networks **Up-Exposure Net** (\mathcal{N}_2) and **Down-Exposure**

Net (\mathcal{N}_3) respectively to hallucinate and generate the high- and low-exposure image with respect to the input image following the Eq. (2).

Masked Regions. In our model, we use masks to eliminate the over- and under-exposed regions in the input image, as Santos *et al.* [36] suggested in their approach. We use masks in both the training and inference phase. Please refer to the supplementary for more details and illustration about the mask generation process.

3.3. Loss function

The proposed network can be trained in an end-to-end fashion. In our training, we sample a pair of low- and high-exposure (I_1, I_2) as input, respectively. By applying the \mathcal{N}_1 network to recover the latent sensor irradiance (\hat{X}_1, \hat{X}_2) , we then scale them by a factor $\frac{\Delta t_i}{\Delta t_j}$ like in Eq. (2) before feeding the outputs to the \mathcal{N}_2 and \mathcal{N}_3 to get the predicted (\hat{I}_2, \hat{I}_1) exposure.

The choice of appropriate loss functions is critical for the learning process of one’s model. In order to synthesize realistic bracketed images, apart from typically used loss functions, we also introduce prior knowledge of the camera formation pipeline to our proposed network. Specifically, we train our network by optimizing the combination of HDR Representation loss \mathcal{L}_h , reconstruction loss \mathcal{L}_r , VGG perceptual loss \mathcal{L}_p and total variation (TV) loss \mathcal{L}_{tv} . Mathematically, the final combined loss function \mathcal{L} is:

$$\mathcal{L} = \lambda_h \mathcal{L}_h + \lambda_r \mathcal{L}_r + \lambda_p \mathcal{L}_p + \lambda_{tv} \mathcal{L}_{tv} \quad (3)$$

HDR Representation Loss. With the setup for our training as shown in Fig. 3 and the image formation pipeline in Fig. 2, we know that (I_1, I_2) comes from the same scene irradiance E with an identical CRF f . Mathematically, we have:

$$\begin{cases} I_1 &= f(X_1) = f(E\Delta t_1) \\ I_2 &= f(X_2) = f(E\Delta t_2) \end{cases} \quad (4)$$

We could infer that, if the inverse CRF f^{-1} is known, I_1 and I_2 only differ by scalars Δt_1 and Δt_2 respectively. Our method relies on this prior knowledge to constrain \mathcal{N}_1 . We introduce the *transformation loss* when transforming from \hat{X}_1 to \hat{X}_2 , following the Eq. (2), as:

$$\mathcal{L}_t(\hat{X}_1, \hat{X}_2) = \left\| \log\left(\hat{X}_1 \odot \frac{\Delta t_2}{\Delta t_1} + \epsilon\right) - \log(\hat{X}_2 + \epsilon) \right\|_1 \quad (5)$$

where Δt_i is the exposure time of the corresponding \hat{X}_i and ϵ is a small constant to prevent numerical error. We take the logarithm of both encoded sensor irradiance before computing loss as computing the loss in the log domain reduces the influence of these substantial errors and encourages the network to restore more details in other regions. Then the *HDR Representation loss* is defined as:

$$\mathcal{L}_h = \mathcal{L}_t(\hat{X}_1, \hat{X}_2) + \mathcal{L}_t(\hat{X}_2, \hat{X}_1) \quad (6)$$

with \hat{X}_1, \hat{X}_2 is the output from \mathcal{N}_1 given the input is I_1, I_2 respectively. The \mathcal{L}_h loss can be seen as a relaxed version instead of directly forcing \hat{X}_1, \hat{X}_2 , e.g., $\left\| \frac{\hat{X}_1}{\Delta t_1} - \frac{\hat{X}_2}{\Delta t_2} \right\|_1$. If our network can learn to predict $\hat{X}_i = X_i$, the loss function would reach its minimum. The reason for using this loss function is that directly constraining \hat{X}_i by multiplying or dividing could lead to exploding or vanishing gradient since Δt usually lies within $[0, 1]$. This would make the training very unstable. Therefore, the loss function in Eq. (6) that we have derived is more suitable for training.

Reconstruction Loss. For supervising up and down-exposure networks (\mathcal{N}_2 and \mathcal{N}_3), this can be seen as an image-to-image translation task in which the typical losses used are pixel-wise ℓ_1 -norm and ℓ_2 -norm. Previous work [41, 46] has shown that ℓ_1 is more effective for preserving details, thus we employ our *reconstruction loss* as:

$$\mathcal{L}_r = \left\| \hat{I}_1 - I_1 \right\|_1 + \left\| \hat{I}_2 - I_2 \right\|_1 \quad (7)$$

where \hat{I}_2, \hat{I}_1 is the output from $\mathcal{N}_2, \mathcal{N}_3$ given the input image is I_1 and I_2 respectively.

Perceptual Loss. The *perceptual loss* is used to evaluate how well the extracted features from prediction are matched

with the extracted features from the ground truth. This can help alleviate artifacts and create more realistic details in inferred images. Following setup in [21, 36] we define our *perceptual loss* as follows:

$$\mathcal{L}_p = \sum_l \left\| \phi_l(\hat{I}_1) - \phi_l(I_1) \right\|_1 + \sum_l \left\| \phi_l(\hat{I}_2) - \phi_l(I_2) \right\|_1 \quad (8)$$

where $\phi_l(\cdot)$ is the extracted features at the l^{th} layer of VGG network. In this work, we use VGG-19 network [39] and feature vectors are extracted from *pool1, pool2, pool3* layers. **Total Variation Loss** To avoid overfitting when training and improve spatial smoothness of inferred images, we also minimize the *total variation loss*. The total variation [35] for an image y can be expressed as:

$$\mathbf{V}(y) = \sum_{i,j} \left(\|y_{i+1,j} - y_{i,j}\|_1 + \|y_{i,j+1} - y_{i,j}\|_1 \right) \quad (9)$$

where i, j are the corresponding pixel coordinates of that image. Given the above definition, our *total variation loss* is calculated on inferred images as:

$$\mathcal{L}_{tv} = \mathbf{V}(\hat{I}_1) + \mathbf{V}(\hat{I}_2) \quad (10)$$

3.4. Inference process

To perform inference, we let the network take a *single LDR image* as input and produce multiple up- and down-exposures of the input image to synthesize an image bracket. Specifically, different exposures are synthesized by scaling the latent scene irradiance of the input with an exposure ratio as discussed in Sec. 3.2. We can then aggregate the images in the bracket to form an HDR image by following the conventional HDR imaging pipeline [4, 23]. In this work, we use Photomatix [11] to generate the HDR image. Tone mapping is then followed to display the HDR image. More details of the inference process are provided in the supplementary.

4. Experimental Results

4.1. Implementation details

Dataset. We used the dataset synthesized by Endo *et al.* [6] for training and testing. The dataset was created by applying five representative CRFs from Grossberg and Nayar's Database of Response Functions (DoRF) [9] on 1,043 collected HDR images with nine exposure values. The process results in a total of 46,935 LDR images for training and 6,210 images for testing. Each image has a size of 512×512 . Images in the dataset cover a wide variety of scenarios such as indoor, outdoor, night, and day scenes. Since our training pipeline receives only pairs of images, we randomly sample two images from each scene in the training set and use them as input to train our model. We do not evaluate on




Figure 4: Tone-mapped HDR images comparison between different methods. DrTMO [6] and Deep Recursive HDRI [18] produce artifacts in extremely high dynamic range regions, SingleHDR [21] appears to have checkboard artifacts, while our method can recover details in these regions pleasingly.

Method	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)	HDR-VDP-2 (\uparrow)
Lee <i>et al.</i> [18]	19.56	0.7920	0.2096	53.86 \pm 4.46
Endo <i>et al.</i> [6]	21.60	0.8493	0.1592	54.56 \pm 4.29
Liu <i>et al.</i> [21]	19.77	0.7832	0.2001	52.77 \pm 5.40
Liu <i>et al.</i> [21]*	21.58	0.8333	0.1427	56.42 \pm 4.50
Ours	23.74	0.8916	0.1231	55.69 \pm 5.01

Table 2: Quantitative comparisons on HDR images. Red and blue text indicates the best and second-best respectively. * indicates that the model has been pretrained on HDR-SYNTH [21].

other datasets such as HDR-SYNTH, HDR-REAL from Liu *et al.* [21], and RAISE [3] because they do not include image pairs with known exposures that are required for our training. Specifically, we investigated the HDR-SYNTH, HDR-REAL data and found that their multiple exposure images are not well organized. The exposures do not match the index on file names, and the images do not include any EXIF tags that can be used for recovering the exposure information.

Training details. Our model is trained using Adam optimizer [15] with batch size and learning rate as 64 and

1×10^{-4} , respectively. We decrease the learning rate by a factor of 0.5 every time the loss reaches a plateau. For each image in the input pair, we randomly crop a patch of 256×256 from it. The cropped image also gets randomly rotated, shifted, scaled, and flipped horizontally and vertically for augmentation, which enriches the input data and prevents the model from overfitting. We implement our model using PyTorch [28] and train it on $2 \times$ NVIDIA Tesla A100 GPU with approximate 200,000 steps for our model to converge. The training phase took about three days to complete.

4.2. Evaluation results

Evaluation protocol. To demonstrate our model’s ability to generate realistic images, we conduct experiments to compare our method against Endo *et al.* (DrTMO) [6], Lee *et al.* [18] (DRHDRI), and Liu *et al.* [21]. We also considered Deep Chain HDRI [17] but could not compare due to the lack of publicly available implementation. The model proposed by Lee *et al.* [18] and their evaluation protocol only use five images with EV ranging from -2 to +2 given the

Method	Reinhard <i>et al.</i> [33]			Photomatix [11]		
	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
Lee <i>et al.</i> [18]	22.68	0.9017	0.1438	21.42	0.8730	0.1785
Endo <i>et al.</i> [6]	27.19	0.9419	0.1289	22.20	0.8944	0.1642
Liu <i>et al.</i> [21]	23.22	0.8954	0.1400	19.90	0.8548	0.1651
Liu <i>et al.</i> [21]*	26.19	0.9135	0.0969	23.89	0.8856	0.1236
Ours	29.68	0.9586	0.0617	25.22	0.9370	0.0933

Table 3: Quantitative comparisons on tone-mapped images with existing methods. **Red** and **blue** text indicates the best and second-best respectively. * represents that the model has been pretrained on HDR-SYNTH [21]. The proposed method outperforms all the others.

Component		HDR				Reinhard Tonemap			Photomatix Tonemap		
Two Exposure Nets	HDR Loss	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)	HDR-VDP-2 (\uparrow)	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
\times	\times	17.51	0.5068	0.3169	52.41 \pm 5.54	15.99	0.7764	0.2421	19.04	0.8009	0.2585
\times	\checkmark	14.48	0.4068	0.3464	50.98 \pm 5.43	14.41	0.7571	0.2107	17.10	0.7804	0.2359
\checkmark	\times	23.05	0.8868	0.1192	56.68 \pm 4.81	29.29	0.9608	0.0580	25.19	0.9396	0.0873
\checkmark	\checkmark	23.74	0.8916	0.1231	55.69 \pm 5.01	29.68	0.9586	0.0617	25.22	0.9370	0.0933

Table 4: Ablation study. Using separate exposure networks or HDR representation loss alone leads to degraded performance.

input image has the EV of 0, each image is different by 1 EV, to reconstruct the HDR image. Thus, we decided to use the same setup as them. For Endo *et al.* [6], the model produces a total of 16 images, with eight images for up and the rest for down-exposed images. We decided to select only five images with the EV difference in the range of -2 to +2 from these for a fair evaluation between different models. For Liu *et al.* [21], we include two versions: with and without pretraining on HDR-SYNTH dataset [21]. Note that Liu *et al.*'s method requires ground truth HDR for training.

Comparisons on HDR images. We use Photomatix [11] to recover the final HDR image from the predicted exposures. The peak signal-to-noise ratio (PSNR), structure similarity (SSIM), LPIPS [45], and HDR-VDP-2 [26] metrics are used to evaluate our reconstructed HDR from bracketed images. The result is shown in Tab. 2. Our model outperforms all previous works in PSNR, SSIM, and LPIPS. Compared to Liu *et al.* [21], our model outperforms in the HDR-VDP-2 metric if their model is not pretrained on HDR-SYNTH. Their pretrained version is slightly better in HDR-VDP-2 metric due to the use of extra training data.

Comparisons on tone-mapped images. The tone mapping operator (TMO) that we use to map the reconstructed HDR image into displayable LDR one is Reinhard *et al.*'s method [33], a popular global TMO that models the human visual system. We also consider the TMO from Photomatix [11] to validate the consistency in quantitative results between different methods. Table 3 shows that our proposed method outperforms all of the others in terms of all metrics - PSNR, SSIM, and LPIPS.

Figure 4 shows our tone-mapped HDR images along with others. DrTMO [6] and Deep Recursive HDRI [18] produce

Masked over-/under-exposure	PSNR (\uparrow)	SSIM (\uparrow)	LPIPS (\downarrow)
\times	16.62	0.8423	0.1681
\checkmark	16.67	0.8581	0.1391

Table 5: Analysis on HDR Encoding Net (\mathcal{N}_1).

artifacts in extremely high dynamic range regions. While SingleHDR [21] could handle these regions, the results appear to have checkboard artifacts. Our method can recover these regions with reasonable details and without artifacts. More qualitative results can be found in the supplementary material.

Figure 5 compares four different exposure values synthesized by Deep Recursive HDRI [18], DrTMO [6], and ours, along with the corresponding ground truth. In the lowest EV, our model can synthesize details in the blue sky that is most plausible and near the ground truth image than the other two methods as DrTMO [6] seems to suffer from artifacts when trying to fill in the details. While in the up-exposure scenario, all considered methods perform reasonably well, but the color and contrast of Deep Recursive HDRI [18] do not match the ground truth images. Their network does not seem to model the CRF well enough to preserve the specific non-linear mapping in the CRF.

Ablation study. We provided an ablation study in Tab. 4 to highlight the effectiveness of our proposed components. Similar to the previous works [6, 18, 19, 19], we empirically found that using a single network to learn up-/down-exposure is ineffective. This study also confirms the benefit of our proposed HDR loss for the PSNR metric.

Network structure analysis. To understand our proposed network, we provide an analysis of the effect of our HDR

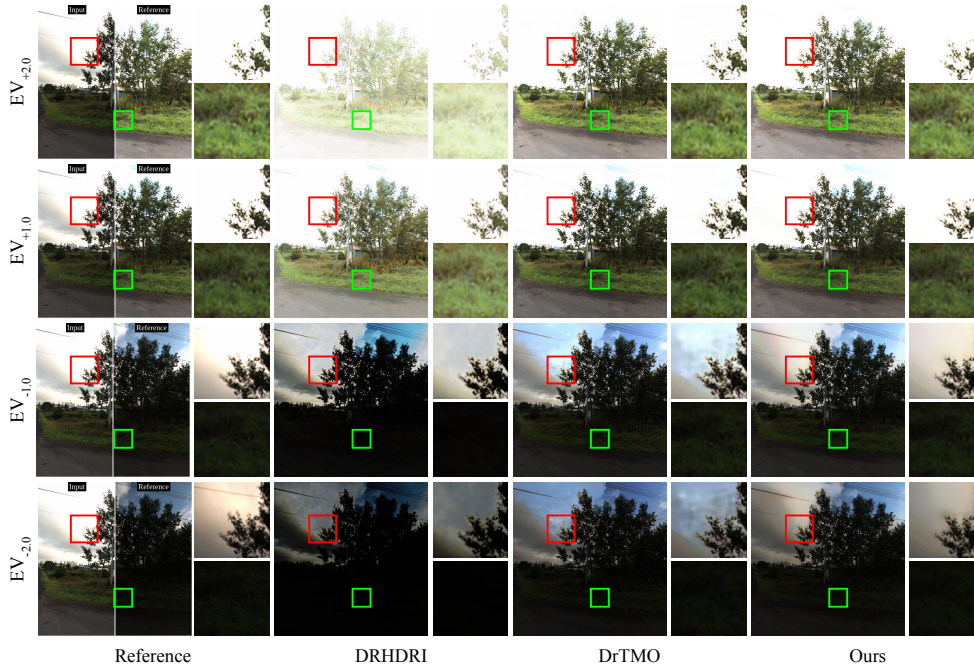


Figure 5: Comparison of bracketed images generated by our model, DrTMO [6], Deep Recursive HDRI [18], and the reference. The overall structure of images is well reconstructed as well as perceptually similar to ground truth images for our method.




Figure 6: Results of user study - Ours vs. the Others. The proposed method is preferred to other methods in both tests.

Encoding Net (network \mathcal{N}_1 in Fig. 3). We take an LDR image I as input to \mathcal{N}_1 . Each LDR image in the dataset is associated with a ground truth camera response function g , which we use to apply to the output of \mathcal{N}_1 . We evaluate the similarity of I and $g(\mathcal{N}_1(I))$ over all images in the test set using PSNR, SSIM, and LPIPS. The results are shown in Tab. 5. It can be seen that \mathcal{N}_1 recovers the scene radiance to a certain extent, reflected by a mid-range PSNR, SSIM, and LPIPS. Visualizations of the output of \mathcal{N}_1 can be found in the supplementary material.

4.3. User study

We conduct a user study on 40 samples to evaluate human preference on the qualitative results. We randomly pick 28 scenes from the total of 33 scenes and show each in pairs [16, 34]. The participants are asked to pick a better image in each pair. First, in the test without reference, we show only two images, one is ours, and one is from the other method to evaluate *visual quality* of the two HDR

reconstructions. Then we do a reference test, in which we add an input LDR and a ground truth HDR image to each question. This is to evaluate the *faithfulness* of each method. We report the detailed comparison in Fig. 6, which shows that we are preferred in both tests compared to all other methods by approximately 70% of the users.

5. Conclusion

We proposed a method for predicting multiple exposures from a single input image, which is applied for HDR image reconstruction. Our method contributes toward making data-driven HDR reconstruction more accessible without the need for ground-truth HDR images for training. Our method can generate an arbitrary number of different exposure images, enabling more granular control in the HDR reconstruction. We achieved state-of-the-art performance on DrTMO dataset [6].

A limitation of our method is that our reconstruction might have visual artifacts or missing details, which we hypothesize this is due to the diversity of the DrTMO dataset as it includes both natural outdoor and man-made indoor scenes, making the hallucination extremely challenging to learn. Future research might integrate generative modeling to improve the image quality or condition the reconstruction on example photographs.

Acknowledgment. This work is done when Quynh Le was a resident of the AI Residency program at VinAI Research.

References

- [1] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11036–11045, 2019.
- [2] William J Claff. Photographic dynamic range versus iso setting, 03 2021.
- [3] Duc-Tien Dang-Nguyen, Cecilia Pasquini, Valentina Conotter, and Giulia Boato. Raise: A raw images dataset for digital image forensics. In *Proceedings of the 6th ACM multimedia systems conference*, pages 219–224, 2015.
- [4] 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, SIGGRAPH '97*, page 369–378, USA, 1997. ACM Press/Addison-Wesley Publishing Co.
- [5] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafal K. Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *ACM Trans. Graph.*, 36(6), nov 2017.
- [6] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Trans. Graph.*, 36(6), nov 2017.
- [7] James A Ferwerda, Sumanta N Pattanaik, Peter Shirley, and Donald P Greenberg. A model of visual adaptation for realistic image synthesis. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 249–258, 1996.
- [8] M.D. Grossberg and S.K. Nayar. Determining the camera response from images: what is knowable? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(11):1455–1467, 2003.
- [9] M.D. Grossberg and S.K. Nayar. What is the space of camera response functions? In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–602, 2003.
- [10] Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy. Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Trans. Graph.*, 35(6), nov 2016.
- [11] HDRsoft. Photomatix.
- [12] Jun Hu, Orazio Gallo, Kari Pulli, and Xiaobai Sun. Hdr deghosting: How to deal with saturation? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1163–1170, 2013.
- [13] Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *ACM Trans. Graph.*, 36(4), jul 2017.
- [14] Zeeshan Khan, Mukul Khanna, and Shanmuganathan Raman. Fhdr: Hdr image reconstruction from a single ldr image using feedback network. In *2019 IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1–5, 2019.
- [15] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.
- [16] Wei-Sheng Lai, Jia-Bin Huang, Zhe Hu, Narendra Ahuja, and Ming-Hsuan Yang. A comparative study for single image blind deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1701–1709, 2016.
- [17] Siyeong Lee, Gwon Hwan An, and Suk-Ju Kang. Deep chain hdri: Reconstructing a high dynamic range image from a single low dynamic range image. *IEEE Access*, 6:49913–49924, 2018.
- [18] Siyeong Lee, Gwon Hwan An, and Suk-Ju Kang. Deep recursive hdri: Inverse tone mapping using generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 596–611, 2018.
- [19] Siyeong Lee, So Yeon Jo, Gwon Hwan An, and Suk-Ju Kang. Learning to generate multi-exposure stacks with cycle consistency for high dynamic range imaging. *IEEE Transactions on Multimedia*, 23:2561–2574, 2021.
- [20] Ru Li, Chuan Wang, Shuaicheng Liu, Jue Wang, Guanghui Liu, and Bing Zeng. Uphdr-gan: Generative adversarial network for high dynamic range imaging with unpaired data. *arXiv preprint arXiv:2102.01850*, 2021.
- [21] 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 (CVPR)*, pages 1651–1660, 2020.
- [22] D. Marnerides, T. Bashford-Rogers, J. Hatchett, and K. De-battista. Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. *Computer Graphics Forum*, 37(2):37–49, 2018.
- [23] T. Mertens, J. Kautz, and F. Van Reeth. Exposure fusion: A simple and practical alternative to high dynamic range photography. *Computer Graphics Forum*, 28(1):161–171, 2009.
- [24] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784*, 2014.
- [25] Kenta Moriwaki, Ryota Yoshihashi, Rei Kawakami, Shaodi You, and Takeshi Naemura. Hybrid loss for learning single-image-based hdr reconstruction. *arXiv preprint arXiv:1812.07134*, 2018.
- [26] Manish Narwaria, Rafal Mantiuk, Mattheiu P Da Silva, and Patrick Le Callet. Hdr-vdp-2.2: a calibrated method for objective quality prediction of high-dynamic range and standard images. *Journal of Electronic Imaging*, 24(1):010501, 2015.
- [27] Yuzhen Niu, Jianbin Wu, Wenxi Liu, Wenzhong Guo, and Rynson WH Lau. Hdr-gan: Hdr image reconstruction from multi-exposed ldr images with large motions. *IEEE Transactions on Image Processing*, 30:3885–3896, 2021.
- [28] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E.

- Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [29] 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 *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 497–513. Springer, 2020.
- [30] K Ram Prabhakar, Rajat Arora, Adhitya Swaminathan, Kunal Pratap Singh, and R Venkatesh Babu. A fast, scalable, and reliable deghosting method for extreme exposure fusion. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–8. IEEE, 2019.
- [31] 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 (ACCV)*, November 2020.
- [32] K Ram Prabhakar, V Sai Srikar, and R Venkatesh Babu. Deep-fuse: A deep unsupervised approach for exposure fusion with extreme exposure image pairs. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 4714–4722, 2017.
- [33] E. Reinhard and K. Devlin. Dynamic range reduction inspired by photoreceptor physiology. *IEEE Transactions on Visualization and Computer Graphics*, 11(1):13–24, 2005.
- [34] Michael Rubinstein, Diego Gutierrez, Olga Sorkine, and Ariel Shamir. A comparative study of image retargeting. In *ACM SIGGRAPH Asia 2010 Papers, SIGGRAPH ASIA '10*, New York, NY, USA, 2010. Association for Computing Machinery.
- [35] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992.
- [36] Marcel Santana Santos, Tsang Ing Ren, and Nima Khademi Kalantari. Single image hdr reconstruction using a cnn with masked features and perceptual loss. *ACM Trans. Graph.*, 39(4), jul 2020.
- [37] 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), nov 2012.
- [38] Ana Serrano, Felix Heide, Diego Gutierrez, Gordon Wetstein, and Belen Masia. Convolutional sparse coding for high dynamic range imaging. In *Proceedings of the 37th Annual Conference of the European Association for Computer Graphics, EG '16*, page 153–163, Goslar, DEU, 2016. Eurographics Association.
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [40] 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.
- [41] 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 (CVPR)*, pages 1751–1760, 2019.
- [42] 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.
- [43] Xin Yang, Ke Xu, Yibing Song, Qiang Zhang, Xiaopeng Wei, and Rynson WH Lau. Image correction via deep reciprocating hdr transformation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1798–1807, 2018.
- [44] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Cycleisp: Real image restoration via improved data synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2696–2705, 2020.
- [45] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 586–595, 2018.
- [46] Hang Zhao, Orazio Gallo, Iuri Frosio, and Jan Kautz. Loss functions for image restoration with neural networks. *IEEE Transactions on Computational Imaging*, 3(1):47–57, 2017.