

Luminance Attentive Networks for HDR Image and Panorama Reconstruction

Hanning Yu^{1*}, Wentao Liu^{1*}, Chengjiang Long^{2†}, Bo Dong³, Qin Zou¹, Chunxia Xiao^{1†‡}

¹ School of Computer Science, Wuhan University, Wuhan, Hubei, China 430072

² JD Finance America Corporation, Mountain View, CA, USA 94043

³ SRI international, Princeton, NJ, USA 08540

fishaning@whu.edu.cn, liu136583410@gmail.com, cjfykx@gmail.com, dongshuhao12@gmail.com, qzou@whu.edu.cn, cxxiao@whu.edu.cn

Abstract

It is very challenging to reconstruct a high dynamic range (HDR) from a low dynamic range (LDR) image as an ill-posed problem. This paper proposes a luminance attentive network named LANet for HDR reconstruction from a single LDR image. Our method is based on two fundamental observations: (1) HDR images stored in relative luminance are scale-invariant, which means the HDR images will hold the same information when multiplied by any positive real number. Based on this observation, we propose a novel normalization method called "HDR calibration" for HDR images stored in relative luminance, calibrating HDR images into a similar luminance scale according to the LDR images. (2) The main difference between HDR images and LDR images is in under-/over-exposed areas, especially those highlighted. Following this observation, we propose a luminance attention module with a two-stream structure for LANet to pay more attention to the under-/over-exposed areas. In addition, we propose an extended network called panoLANet for HDR panorama reconstruction from an LDR panorama and build a dual-net structure for panoLANet to solve the distortion problem caused by the equirectangular panorama. Extensive experiments show that our proposed approach LANet can reconstruct visually convincing HDR images and demonstrate its superiority over state-of-the-art approaches in terms of all metrics in inverse tone mapping. The image-based lighting application with our proposed panoLANet also demonstrates that our method can simulate natural scene lighting using only LDR panorama. Our source code is available at <https://github.com/LWT3437/LANet>.

CCS Concepts

• **Computing methodologies** → **Computer graphics; Artificial intelligence;**

1. Introduction

With a limited dynamic range, a low dynamic range (LDR) image captured by a standard digital camera can not represent the real luminance of the scene and suffers under-/over-exposure problems. High dynamic range (HDR) images solve this problem by recording HDR information of the scene. As HDR images can provide larger luminance variance and contain richer details than LDR images, people can benefit significantly from HDR images in visual perception. HDR images have also been widely used in image-based lighting (IBL) technology to provide more realistic rendering results. A common way to construct an HDR image is by merging a stack of bracketed exposure LDR images. However, dynamic scenes require significant efforts to produce equal-quality results like static scenes, as the moving parts require special treatments.

Therefore, generating an HDR image from a single LDR image is getting more attention. However, reconstructing a high-quality HDR image from a single LDR image is also very challenging.

Most previous methods leverage inverse tone-mapping to solve the problem through traditional image processing technology [AFR*07, KO14, MAF*09, MSG17, RTS*07]. These methods exploit individual heuristics or manual intervention to enhance LDR images, which cannot sufficiently compensate for lost data caused by color quantization and under-/over-exposure. Recently, deep learning-based approaches [EKD*17, EKM17, KTC12, MBRHD18, LHAK18, WZW19] are proposed to automatically infer a statistically plausible HDR image from a single input LDR image. Although these deep learning-based approaches can bring more appealing results, they still have room to be improved. More importantly, most existing HDR datasets are in the relative luminance domain instead of the absolute one, bringing scale ambiguity to the training process. Previous methods used the general maximum normalization method to process HDR images, but this did not solve the scale ambiguity of luminance. This motivates us

† This work was co-supervised by Chengjiang Long and Chunxia Xiao.

* Hanning Yu and Wentao Liu are joint first authors.

‡ Chunxia Xiao is the corresponding author.

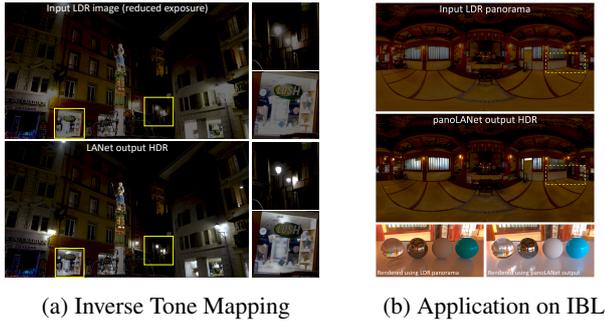


Figure 1: Example results of our proposed method. The left is the reconstruction result of our proposed LANet on a single image. The right results are rendered by original LDR panorama and reconstructed HDR panorama using our panoLANet, respectively. Note that the images of subfigure (a) are mapped using LuminanceHDR software in the same dynamic range and we reduced the exposure of original result for visualizing.

to investigate how to remove the uncertainty caused by the different luminance scales and achieve better reconstruction results.

We observe that two different HDR images may store the same scene, as they are saved in different luminance scales. That is to say, the difference between two HDR images is scale-invariant in the relative luminance domain, which means multiplying an HDR image by any positive real number will not change the information it represents. In such cases, to deal with the scale ambiguity problem, we propose an HDR normalizing approach to calibrate all HDR images into a similar luminance scale and introduce the scale-invariant loss [EPF14] into the HDR learning task.

We also observe that the main error between LDR images and HDR images occurs in the under-/over-exposed areas, especially in over-exposed areas. This requires the network to reconstruct the high dynamic range based on the image's brightness. Here we introduce the attention mechanism with two different network structures to achieve local spatial attention for general images and panoramas respectively, enabling the networks to focus more on over-exposed areas while avoiding errors to other areas as much as possible.

In this paper, based on the above two observations, we propose Luminance Attentive Network (LANet) to construct an HDR image from a single LDR image, as shown in Fig. 2. The LANet is designed as a multi-task network with two streams, named *luminance attention stream* and *HDR reconstruction stream*. The luminance attention stream is designed for network to learn to obtain a spatial weighted attention map about the luminance distribution. This design exploits estimated luminance segmentation as an auxiliary task to supervise the attention weights, and a novel luminance attention module is proposed to guide the reconstruction process paying more attention to those under-/over-exposed areas.

To validate the potential of our proposed LANet, we further propose its extension network called panoLANet (see Fig. 5) to reconstruct HDR panoramas for IBL rendering. The panoLANet consists of a *ceiling luminance branch* and a *panorama reconstruction branch*, in which the ceiling luminance branch is used to reconstruct the highlighted area in the upper half of the scene. We rotate

and apply a perspective transformation on the panorama to produce a ceiling-view image to solve the distortion problem of the equirectangular panorama and use skip-connection with gated attention to sharing the information between the two branches.

It is worth mentioning here that although the usage of HDR panoramas is different from general HDR images, we still design these two networks, LANet and panoLANet, with similar structures and design principles. Both of these two networks can use the same data processing and training methods. LANet, as a general inverse tone mapping method, can achieve the HDR reconstruction task from the most common images. Meanwhile, as an extension of LANet, panoLANet can be applied to IBL rendering technology. As two examples, Fig. 1 (a) shows our proposed LANet can well handle the areas under-/over-exposed areas, and Fig. 1 (b) shows that our proposed panoLANet can generate high-quality results for applications on image-based lighting.

To sum up, the main contributions of this paper are three-fold:

- We propose an end-to-end trainable luminance attentive network called LANet with two streams for HDR reconstruction and performs better quantitatively and qualitatively than prior work.
- We propose a novel HDR calibration method, calibrating HDR images into a similar luminance scale according to the LDR images.
- We extend LANet to the panoLANet for HDR reconstruction from an LDR panorama.

We conduct experiments on public datasets [NKHE15, Hdr12] with both real LDR inputs and ground truth HDR references. Extensive experimental results have powerfully demonstrated the superiority of our proposed approaches LANet and panoLANet over state-of-the-art approaches.

2. Related Work

The related work can be divided into three categories, *i.e.*, *custom-made hardware*, *inverse Tone-Mapping Operators (iTMO)*, and *Deep learning-based approaches*.

Custom-made hardware is one way to generate high-quality HDR images from a single shot through specially designed hardware (*e.g.*, coded aperture and beam-splitter) [TKTS11, MMP*07, Pul14, HKU15]. However, custom-made hardware is more expensive and high-profile, making it hard to use widely.

Inverse tone-mapping refers to reconstructing HDR from a single LDR image. The earlier approaches [BLDC06, MAF*09, AFR*07, MAF*09] tried to achieve this goal by applying an expand function to the LDR image. Filter based methods were also proposed to deal with both images and videos [RTS*07, KO14]. However, these approaches are not user-friendly since they requires to adjust various parameters, which is troublesome for non-expert users to achieve desired results.

Deep learning-based approaches have made remarkable achievements in many visual applications like visual recognition [LH15, LH17, HLYG18, HLX21b, HLX21a], object detection [LWH*14, ILBH20], super resolution [ZLW*21, ZLW*20], image captioning [DLXX21], shadow detection and removal [DLZX19, WLZX19, ZLZX20, ZLY*20, CLZX21], shadow

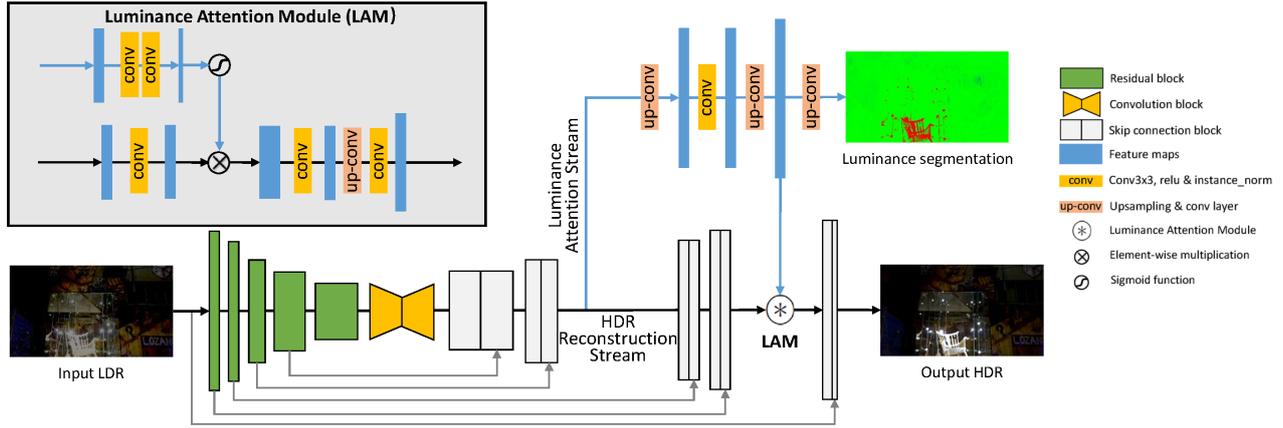


Figure 2: The overview of our LANet network. The convolution block consists of two convolutional layers with stride 2 and two up-conv layers. The skip-connection block first upsamples the previous layer’s output and through two 3×3 convolution layers to keep the number of channels consistent with the jump layer. Meanwhile, a 3×3 convolution layer is applied to the output of the skip layer. Finally, the two convolutional results are concatenated and through an 1×1 convolution layer to obtain the output result.

generation [ZLW19, LLZ*20], action localization [ILR21], trajectory prediction [SWL*21] and so on. Deep learning also provides another way to generate a high-quality HDR image from a single LDR image. Instead of directly producing HDR images by a DCNN model, Endo *et al.* [EKM17] developed a DCNN model to generate a stack of bracketed exposure images from a single LDR image. The same idea was adopted by Lee *et al.* [LAK18, LHAK18], where they used a chaining structure and GAN network to obtain the required bracketed exposure images. However, the HDR images generated by these methods still have a limited dynamic range.

Recently, directly generating HDR images from a DCNN model is getting more attention. Eilertsen *et al.* [EKD*17] developed a virtual camera to generate LDR from HDR datasets, providing enough LDR-HDR pairs to train a model for directly generating HDR image. However, they only predicted the overexposure areas, and the virtual camera was created based on an out-of-date cameras database. Marnierides *et al.* [MBRHD18] achieved the same goal with a novel network structure by splitting the LDR encoder into local, dilation, and global branches. Kim *et al.* [KOK19] learned super-resolution and inverse tone-mapping for UHD HDR applications, but they focus on the final tone-mapped images for display technology and are not concerned with the images in HDR format. Zhang and Aydın [ZA21] decomposed an input LDR image into a base and detail layer and Liu *et al.* [LLC*20] achieved single image HDR reconstruction by learning to reverse the camera pipeline. However, the step-by-step strategy leads to the accumulation of errors among the subnetworks. Zhang *et al.* [ZL17] focused on sun-light outdoor lighting estimation on panoramas, but they need to rotate the panorama in advance so that the sun is in the middle of the image. Unlike the existing deep learning methods, we start from the essential properties of HDR images to pre-process the data by calibrating HDR images into a similar scale and introduce the attention mechanism with scale invariance to build our networks. In this way, we propose corresponding HDR reconstruction approaches for both



Figure 3: An example of HDR training batch with and without our calibration method. These two batched images are mapped using the LumaninceHDR software in a certain exposure and dynamic range for visualizing.

images and panoramas, which can be applied to both HDR display and IBL applications.

3. Luminance Attentive Network

As illustrated in Fig. 2, we propose a novel luminance attentive network named LANet, designed with two streams and fully leverages the auxiliary luminance segmentation for better recovering an HDR image. It has three essential components, *i.e.*, HDR calibration with luminance scale invariance, luminance attention module, and luminance scale-invariant loss. We are going to discuss each component as well as the luminance segmentation learning in the following subsections.

3.1. HDR Calibration with Luminance Scale Invariance

The main problem of storing HDR images in relative luminance is that the luminance scales may differ. If we use them as ground truth for training a DCNN model, the scale ambiguity will confuse

the training process. Therefore, to remove the ambiguity, we propose a novel HDR calibration approach to provide ground truth for training our LANet (see Section 3.3). The key to our calibration is luminance scale invariance. In the relative luminance domain, we call two HDR images H_1 and H_2 luminance scale-invariant if and only if they are only differed by a positive scale factor $\kappa \in \mathcal{R}^+$ such that $H_1 = \kappa H_2$. Our goal is to make all HDR images at the same luminance level as much as possible. Since all LDR images are already in the same luminance units and fewer errors in the non-overexposed areas between LDR and HDR images, it is reasonable to choose the LDR version as the standard level to calibrate all of the HDR images. To find the most effective scale factor, for each HDR image H , we scale it to the same dynamic range with the corresponding normalized LDR image I (i.e., $I(x, y) \in [0, 1]$). The calibrated HDR image \hat{H} can be obtained as:

$$\hat{H} = \frac{S(M \odot I)}{S(M \odot H)} \cdot H \quad (1a)$$

$$S(P) = \sum_{c,y,x} P(x, y, c), \quad (1b)$$

$$M(x, y) = \begin{cases} 1, & \text{if } \frac{1}{3} \sum_c I(x, y, c) < \tau \\ 0, & \text{otherwise} \end{cases} \quad (1c)$$

where \odot denotes element-wise multiplication; $P(x, y, c)$ is a pixel value at position of (x, y) on channel c of an image P ; $\tau \in [0, 1]$ is a threshold used to determine non-overexposed areas, we set $\tau = 0.83$ for our experiments which can achieve the best performance in our test. Note that I is in the linear RGB space.

As we can see, after applying Eq. 1a, the HDR image will contain a similar value with the corresponding LDR image in non-overexposed areas, and only the highlight pixels will have a significant difference. As shown in Fig. 3, we can see that images after alignment look more regular at the same specific, while original data have an extensive bias on the luminance scale. What is more, we can apply thresholding segmentation on the calibrated HDR images to obtain pixel-level luminance segmentation labels automatically without requiring any extra manual annotation efforts.

3.2. Network Architecture and Luminance Attention Module

The proposed LANet is designed as a multi-task learning framework with two output streams, i.e., *luminance attention stream* and *HDR reconstruction stream*. For each LDR image, the LANet predicts the corresponding HDR image and luminance segmentation mask. We adopt U-Net [RFB15] as our backbone structure. Specifically, the encoder follows ResNet50 [HZRS16, HZR16], and each of the five residual blocks connects to the decoder by a skip connection. On the decoder side, the following changes have been made.

- Since a standard deconvolution layer gives checkerboard artifacts [ODO16], we use a nearest-neighbor upsampling and two convolution layers for each up-sample block.
- To better keep the distribution of each HDR image, instance normalization [UVL16] is used in our decoder instead of batch normalization [IS15].
- luminance attention stream branches out from the middle layer of the decoder and follows by four convolution layers and three upsampling to construct a luminance segmentation mask.

- We design a novel luminance attention module (LAM) and incorporate it into the decoder, which brings the luminance attention information from the luminance attention stream to the HDR reconstruction stream.

As shown in Fig. 2, LAM is added before the last skip-connection. It takes the last layer of the luminance attention stream and feature maps generated from the second last skip-connection block as input to construct luminance attention feature maps as input to the last skip-connection. With this treatment, the feature maps from the luminance attention stream are transferred into a two-channel attention map through two convolutional and one sigmoid activation layer. The feature maps from the HDR stream also need to go through one convolutional layer. Then, an element-wise multiplication is applied to each channel of the attention map and the generated feature maps from the HDR stream. The outputs are concatenated in channel dimension as an input to another three convolutional layers. The thoughts behind this design are that the attention scheme should give more guidance for both the training and inferring process.

3.3. Loss Functions

The overall loss function for LANet is formulated with a luminance scale-invariant loss and a luminance segmentation loss, i.e.,

$$\mathcal{L} = \mathcal{L}_{SI} + \alpha \mathcal{L}_{SEG}, \quad (2)$$

where α is a weight hyperparameter to control the trade-off between two losses. Note that $\alpha = 0.05$ is used in our experiments. Inspired by Eilertsen et al. [EKD*17], the proposed LANet predicts HDR \bar{H} in the logarithmic scale, which is better matching how the human visual system reacts to luminance. However, due to \bar{H} is still in the relative luminance domain, dealing with the aforementioned scale ambiguity is required. Therefore, we take advantage of luminance scale invariance and define \mathcal{L}_{SI} as a scale-invariant MSE [GJAF09] in logarithmic scale as:

$$\begin{aligned} \mathcal{L}_{SI}(\bar{H}, \hat{H}) &= \min_{\kappa \in \mathcal{R}^+} \frac{1}{n} \|\log(\kappa \bar{H}) - \log(\hat{H})\|^2 \\ &= \min_{\kappa \in \mathcal{R}^+} \frac{1}{n} \sum_{x,y,c} \left[\log(\bar{H}(x, y, c) + \epsilon) \right. \\ &\quad \left. - \log(\hat{H}(x, y, c) + \epsilon) + \log \kappa \right]^2, \end{aligned} \quad (3)$$

where ϵ is a small constant to avoid $\log 0$ when calculating logarithms. As we can see, Eq. (3) has a closed-form solution, i.e.,

$$\log \kappa = \frac{1}{n} \sum_{x,y,c} (\log(\hat{H}(x, y, c) + \epsilon) - \log(\bar{H}(x, y, c) + \epsilon)).$$

Therefore, with definition $d(x, y, c) = \log \bar{H}(x, y, c) - \log \hat{H}(x, y, c)$, Eq. (3) can be expressed as:

$$\mathcal{L}_{SI}(\bar{H}, \hat{H}) = \frac{1}{n} \sum_{x,y,c} d(x, y, c)^2 - \frac{1}{n^2} \left(\sum_{x,y,c} d(x, y, c) \right)^2. \quad (4)$$

The first component in Eq. (4) indicates the average of square distances, and the second component refers to the square of the mean distance, ensuring the scale-invariance of the loss function. As we can see, the loss value will not change when the outputs plus a single constant in the logarithmic domain. This means that the network treats HDR images in relative luminance.

The segmentation loss \mathcal{L}_{SEG} is a cross entropy loss, *i.e.*,

$$\mathcal{L}_{SEG}(m, \hat{m}) = - \sum_{x,y,c} \left[\hat{m}(x,y,c) \log m(x,y,c) + (1 - \hat{m}(x,y,c)) \log (1 - m(x,y,c)) \right], \quad (5)$$

where m is a predicted luminance mask; \hat{m} is the ground truth luminance mask, which has the same shape as a reconstructed \hat{H} and initialized with 0 for all pixel. Then, we set value as follows:

$$\hat{m}(x,y,c_{idx}) = \begin{cases} 1 \text{ and } c_{idx} = 0, & \text{if } \hat{H}'(x,y) \leq t_l \\ 1 \text{ and } c_{idx} = 1, & \text{if } t_l < \hat{H}'(x,y) < t_h \\ 1 \text{ and } c_{idx} = 2, & \text{if } \hat{H}'(x,y) \geq t_h, \end{cases} \quad (6)$$

$\hat{H}'(x,y)$ is the channel-wise pixel average value at location (x,y) of \hat{H} ; t_l and t_h are the threshold values to define dimmer and brighter area respectively, which are set to $e^{-5.5}$ and $e^{0.1}$ in our experiments.

4. Extension of Luminance Attentive Network for HDR Reconstruction on Panorama

Regarding HDR images for display and rendering, although both require images with real luminance information, there are two significant differences between them. Firstly, the ways to measure the quality of their application are different. HDR display technology displays the HDR image directly to the user for observation, and the effect depends on the perception of the image itself to the human eye. In contrast, the image-based lighting (IBL) technology uses the HDR image as the scene's light source, and the effect depends on the quality of the rendering results. Secondly, the panorama represents a spherical surface in three-dimensional space, which is invariant to spatial rotation. When it is expanded to equirectangular form, a significant distortion will occur in the upper and lower regions of the image.

Given these two differences, we extend our proposed LANet and present a novel two-branch network called panoLANet to improve the HDR reconstruction on panoramas for better rendering results. As illustrated in Fig. 5, our panoLANet consists of a ceiling luminance branch and a panorama reconstruction branch.

Unlike the panorama reconstruction branch, which is used to

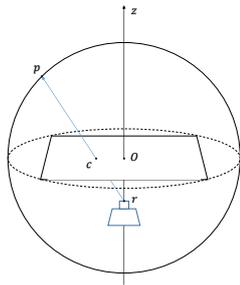


Figure 4: The illustration of the perspective conversion method used in our proposed panoLANet. The entire sphere represents the panorama. The plane at the center of the sphere perpendicular to the z -axis is the ceiling-view image. Position r is the camera center of the perspective projection. For any point p on the upper hemispherical surface of the panorama, the intersection point c of pr and the plane is the corresponding projection point.

roughly predict the entire scene's information, the ceiling luminance branch is designed to accurately reconstruct the highlighted area in the upper half of the scene. We use a differentiable perspective transformation method to convert the panoramas and perspective ceiling-view images (P2C and C2P, respectively). Finally, merge them using the mask to obtain the final reconstructed HDR panorama. In the following subsections, we first describe the method of perspective transformation and then explain the detail of the two-branch network and the loss function.

4.1. P2C and C2P Conversion

The original panorama represents a two-dimensional spherical surface in three-dimensional space, describing the scene around the shooting center. P2C aims to convert most of the information of the upper hemisphere of this two-dimensional sphere into a plane image through perspective projection. Some previous works directly use the center of the sphere as the center of the projection camera to perform the perspective conversion. However, we want to retain the information in the upper half of the space as much as possible. As shown in Fig. 4, we chose to place the center of the projection camera below the center of the sphere.

Given the field of view of the perspective transformation with the resolution of the panorama and ceiling-view image, we can calculate the position of all pixels of the ceiling-view image on panorama or vice versa, shown as the Fig. 4. Using bilinear interpolation to calculate the final value of each pixel, we get the result of P2C and C2P conversion. Fig. 6 shows an example of the conversion of two types of images. It can be seen that the ceiling-view image obtained from P2C can save most of the information of the upper half of the space with less distortion. In addition, by thresholding the ceiling-view image to get the mask of the over-exposed area, we can merge the predicted results of the two branches shown in Fig. 5 to get the final HDR panorama.

4.2. Network and Loss Function

Looking at Fig. 5 with Fig. 2, we can see that panoLANet reuses much of the structure from LANet, which allows us to reduce the training cost through transfer learning. Specifically, the panorama reconstruction branch in panoLANet applies the basic U-Net structure in LANet. The encoder of the ceiling luminance branch also uses the same ResNet50 as the panorama reconstruction branch. We first use the single image data set with an enormous amount of data to train the panorama branch and then use the pretrained model to initialize the panorama reconstruction branch and the encoder of the ceiling luminance branch. Finally, we use the panorama data set to fine-tune the entire panoLANet to get the final training model.

The ceiling luminance branch aims to reconstruct better the HDR of the light source above the space. When we use IBL technology to render, only the upper half of the light affects the rendered object, and the light source is usually above the space rather than below. On the other hand, the light source is the most critical area of the image-based lighting, so the results rendered from LDR panoramas often look unrealistic. We convert the panorama to a ceiling-view image and then build an encoder-decoder network to focus on reconstructing the light source. Specifically, to avoid the loss of in-

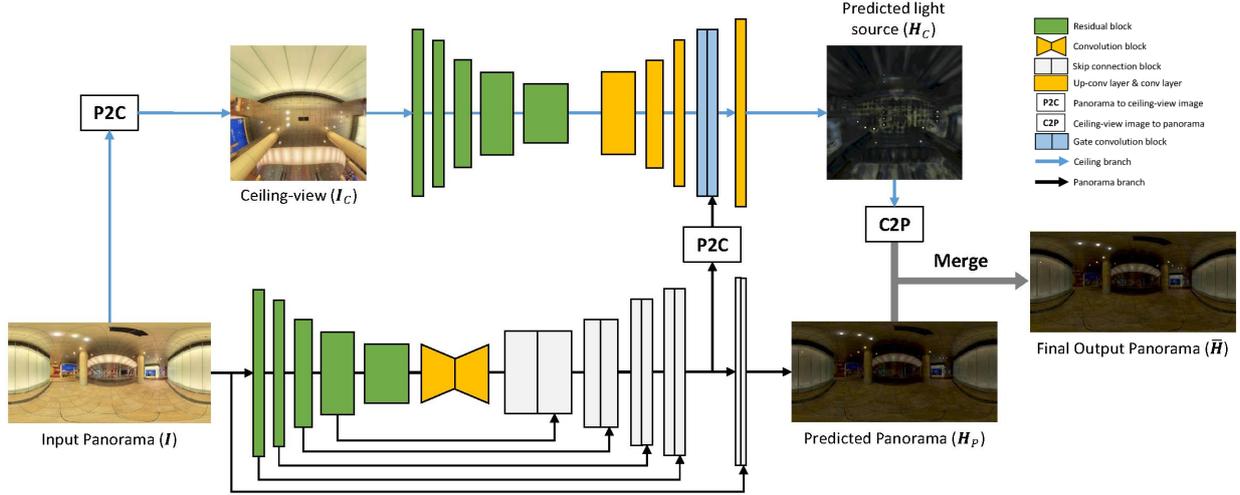


Figure 5: The overview of our panoLANet network, which consists of a ceiling luminance branch and panorama reconstruction branch.



Figure 6: An example result of P2C and C2P conversion: (left) equirectangular panorama, (right) ceiling-view image converted through P2C. Note that when using C2P to convert back to the panorama, the information of the masked area will lose, and therefore we set these uncertain pixels to 0 in our paper.

formation caused by the downsampling operation, we use a skip connection to share the features from the panorama reconstruction branch to the ceiling luminance branch. Based on the feature from skip connection, we use the gated convolution [YLY*19] to achieve local attention, which is called gate convolution block on Fig. 5, so that the network can focus more on the over-exposed areas.

Our goal is to recover HDR information from the LDR panorama for image-based lighting. In order to get the final HDR reconstruction result, we need to merge the prediction results of the two branches. Given the prediction of the ceiling luminance branch H_C and prediction of the panorama branch H_P , the merged final output HDR panorama \bar{H} could be written as

$$\bar{H} = m_p \odot \mathcal{F}_{c2p}(H_C) + (1 - m_p) \odot H_P, \quad (7)$$

where \mathcal{F}_{c2p} is the C2P conversion and \odot denotes element-wise multiplication. The m_p is the mask generated by ceiling-view input I_C as

$$m_p = \frac{\max(0, \text{avg}_c(\mathcal{F}_{c2p}(I_C)) - \tau)}{1 - \tau}, \quad (8)$$

where avg_c means average on the image channels and $\tau = 0.13$ is a threshold to generate the mask. Note that I_C is in linear RGB space.

The loss function in panoLANet is no longer a scale-invariant

loss, because the scale-invariant loss described in Eq. 4 requires that each pixel of the image contributes equally to the loss. Here we use two loss functions \mathcal{L}_{high} and \mathcal{L}_{low} to describe the loss of over-exposed areas and non-overexposed areas, respectively. We define $d = \log \bar{H} - \log \hat{H}$, then the final loss function for panoLANet is

$$\begin{aligned} \mathcal{L}(\bar{H}, \hat{H}) &= \beta_1 \mathcal{L}_{high}(\bar{H}, \hat{H}) + \beta_2 \mathcal{L}_{low}(\bar{H}, \hat{H}) \\ &= \beta_1 |m_p \odot d|^2 + \beta_2 |(1 - m_p) \odot d|^2, \end{aligned} \quad (9)$$

where β_1, β_2 is weight hyperparameters to control the trade-off between two losses. We set $\beta_1 = 0.2, \beta_2 = 0.01$ in our experiments.

5. Experiments

To validate our proposed method, we collect several public HDR datasets shown in Table 1 and conduct various experiments on these datasets. For single image HDR reconstruction of our LANet, we pick out the HDR-Eye dataset [NKHE15] from the collected dataset and use the remaining HDR data to generate training and test data. Since the HDR-Eye dataset has both real LDR and HDR data, we use this dataset as an additional comparison dataset with the previous work to show the applicability of our method. For panorama HDR reconstruction, we use the sIBL [Hdr12] dataset with both LDR and HDR panorama data as our test data.

Table 1: The list of HDR datasets

Type	Dataset Name	Number
Pano	Laval Indoor HDR Dataset [GSY*17]	2233
	Laval Outdoor HDR Dataset [HGAL19]	205
	HDRl Haven [Zaa19]	322
	sIBL [Hdr12]	79
Img	HDR Photographic Survey [Fai17]	105
	Funt et al. HDR Dataset [FS00]	105
	Stanford HDR Data [XDCW02]	88
	Ward [War19]	33
	HDR-Eye [NKHE15]	42
Video	LiU HDRv [KGBU13, KGB*14, DLMM16]	10

5.1. Implementation Detail

We calibrate all HDR images using Eq. 1a and resize the training pairs to the size of 256×256 for general images data and 512×256 for panoramas data. We implement our proposed LANet and panoLANet in TensorFlow and adopt ADAM optimizer [KB15] to train the model with one NVIDIA RTX 2080Ti GPU. The training takes 100K iterations in total with a batch size of 16 for LANet and takes 100K iterations for pretraining with a batch size of 16 and 35k iterations for fine-tuning with a batch size of 4 for panoLANet. Regarding the learning rate, we set the initial learning rate as $4e-5$ and update it with a step decay schedule. Specifically, the learning rate is dropped 0.8 times for every 5000 iterations.

5.2. Datasets and Metrics

The entire HDR raw data we collect contains 2,839 HDR panoramas, 373 single HDR images and 10 HDR videos. For each panorama, we crop images to cover six equidistant azimuths in the horizontal viewing angle and three equidistant azimuths in the 45-degree elevation angle (except outdoor scenes). Each cropped image has an aspect ratio of 4 : 3. We crop images in the center with two aspect ratios for every single HDR image, either 4 : 3 or 3 : 4. For each HDR video, we select the suitable frames to be included in our training data. In this way, we get 25,308 HDR images in total for experiments. Note that most HDR datasets have only HDR images. We need to generate the synthetic LDR images from HDR data. To simulate the modern cameras, we complete our virtual camera by following steps:

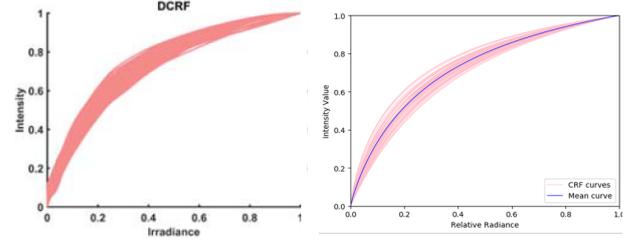
1. We randomly set a dynamic range for the virtual camera from $9.6EV$ to $14.8EV$ based on the popular Digital Single Lens Reflex (DSLR) cameras [DXO19].
2. According to the dynamic range, we use a mean-value auto-exposure algorithm to find an exposure that makes the mean value of the mapped LDR image approximately to middle gray. After that, we obtain the mapped linear LDR image in which all pixels are in the range $[0, 1]$.
3. We apply approximate Camera Response Function (CRF) curves to the linear images, which are in the same form as Eilertsen *et al.* [EKD*17],

$$f(H_{l,c}) = (1 + \sigma) \frac{H_{l,c}^n}{H_{l,c}^n + \sigma}, \quad (10)$$

while we here set the σ in range $[0.3, 0.5]$ and n in range $[0.8, 1.0]$ to fit a modern CRF database. [CMY19]). The modern CRF and our approximate curves are shown in Fig. 7.

The final mapped LDR images are scaled to 0 – 255 in integer to fit the available 8-bits LDR images. For each HDR image, we generate one LDR image using the above virtual camera and the other one with Display Adaptive TMO [MDK08] to fit the images that may be post-processed. We finally got a total of 50616 image pairs in our experiment.

For experiments on general images, we randomly select 90% image pairs for training. For the remaining 10% image pairs, we discard all pairs obtained from Laval HDR Dataset [GSY*17, HGAL19] to avoid testing cases over-concentrated in a majority



(a) Modern CRF curves [CMY19] (b) With our method

Figure 7: The modern CRF curves collected by Chen *et al.* and our random curves. We simulate the CRF curves through a sigmoid function that parameters are set from a certain range.

Table 2: Quantitative Results of Ablation Study on LANet. **N** indicates that the data has been normalized while **C** indicates calibrated.

Method	PU-PSNR \uparrow	PU-SSIM \uparrow	$Q_H\uparrow$	siMSE \downarrow
U-Net+ L_2 (N)	33.09	0.922	38.31	7.93
U-Net+ L_2 (C)	35.04	0.963	41.63	6.70
U-Net+ L_{SI}	35.37	0.966	42.58	6.95
w/o L_{SI} (N)	34.31	0.952	37.33	6.34
w/o L_{SI} (C)	35.24	0.969	40.31	6.30
w/o LAM	35.27	0.965	42.38	7.12
w/o Seg	35.65	0.972	42.59	5.99
w/ LDR-Seg	35.70	0.971	42.33	6.22
w/ HDR-Seg	35.79	0.973	42.53	5.62

dataset and use the rest 802 pairs for evaluation. We also evaluate comparison performance on the HDR-Eye dataset [NKHE15] which contains 46 pairs of ground truth HDR images and LDR images captured from several different cameras. Note that we only have 42 pairs for evaluation because the LDR images have large black regions in the first four pairs.

We remove the data generated from the sIBL dataset for experiments on panoramas and randomly select 90% image pairs for pretraining. Then we use the corresponding panoramas of the pre-trained data to fine-tune the panoLANet model. We evaluate our panoLANet on the sIBL dataset [Hdr12] which contains 79 pairs of ground truth HDR panoramas and corresponding LDR panoramas. Note that we conduct the evaluation on rendering results rather than panorama itself, which will be described in detail in Sec. 5.5.

Regarding the metrics, we use the Q score of HDR-VDP-2 [NMDSL15] as Q_H score, scale-invariant MSE (siMSE) [EPF14], perceptual uniformity encoded PSNR (PU-PSNR) and SSIM (PU-SSIM) [AMS08], to evaluate the quality of output HDR images in HDR domain. Since the HDR-VDP-2, PU-PSNR and PU-SSIM metrics are display-referred metrics applicable for absolute luminance in cd/m^2 , we need to adjust the luminance level to a certain absolute luminance. We first scale the predicted HDR images to the ground truth HDR images, and then we adjust their luminance level approximately to where the maximum value of the corresponding LDR image is $255cd/m^2$. For siMSE metric, we need not change results when comparing this metric because it is completely scale-invariant. As the default

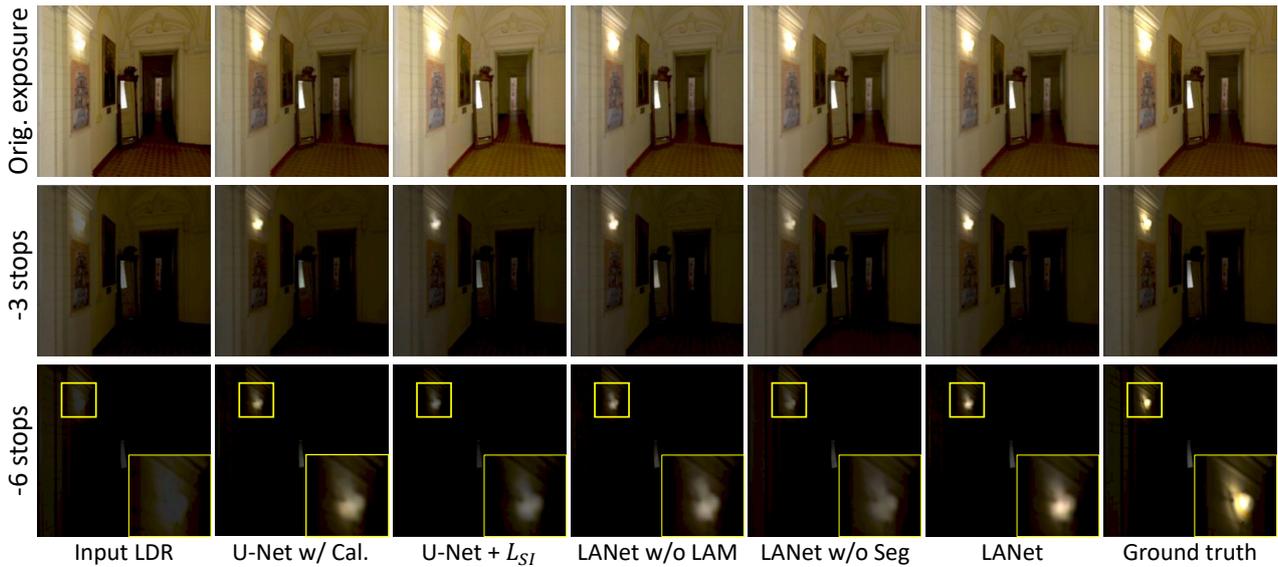


Figure 8: An example of the ablation models. We show the same images in different reduced exposure. For each image, we use the same dynamic range and a manually adjusted exposure with LuminanceHDR software on the original image for the best visual comparison.

configuration for HDR-VDP-2 metric, color encoding is set as “rgb-bt.709” for HDR evaluation, 24-inch display, 1920×1080 resolution and viewing distance of 1 meter.

5.3. Ablation Study on LANet

To evaluate some components of our proposed LANet, we design a series of variants as follows:

- U-Net + L_2 : remove attention stream and LAM and use L_2 loss for training.
- U-Net + L_{SI} : remove attention stream and LAM and use L_{SI} loss for training.
- w/o L_{SI} : use L_2 loss rather than L_{SI} , and without using supervised information for luminance segmentation.
- w/o LAM: remove LAM to make luminance segmentation and HDR reconstruction separately.
- w/o Seg: Full LANet model, but without using supervised information for luminance segmentation.
- w/ LDR-Seg: use the luminance segmentation supervised information calculated based on the input LDR images rather than calibrated HDR images.
- w/ HDR-Seg: use the luminance segmentation supervised information calculated based on the calibrated HDR images.

Meanwhile, we evaluate the general maximum normalization method and our calibration method on “U-Net + L_2 ” and “LANet + L_2 ” which can take HDR data with different luminance scales as ground-truth for training. Note that there is no difference for these two processing when using L_{SI} .

We train the above mentioned variants with our proposed LANet on the same training data and evaluate their results on our self-collected 802 pairs of LDR-HDR images. The results are summarized in Table 2, we use the Q_H score as the value of HDR-VDP-2 metric and the results of siMSE metric are in 10^{-2} units. From

the results, we can observe: (1) for both U-Net+ L_2 and LANet+ L_2 , the performance with calibrated HDR as ground-truth for training is better than taking general normalized HDR as ground-truth; (2) with L_{SI} loss, the performances of U-Net and LANet are better than with L_2 loss, which verifies the effectiveness of scale invariance in HDR; (3) with L_{SI} loss, all of the LANet model performs better than U-Net, which suggests that LAM works with or without using luminance segmentation labels for supervision; (4) LANet shows superiority when compared with LANet w/o LAM, which demonstrates the efficacy of LAM with luminance attention stream; (5) with LDR-Seg to replace HDR-Seg, LANet w/ LDR-Seg achieves lower performance than LANet, which indirectly demonstrate the effectiveness of luminance segmentation labels obtained from calibrated HDR rather than input LDR images; and (6) LANet w/ HDR-Seg achieves the best performance in all the metrics except Q_H score, which shows that using luminance segmentation supervision can improve the performance. Because adding the segmentation supervision does not require additional calculations, we use the LANet w/ HDR-Seg as the final model of LANet.

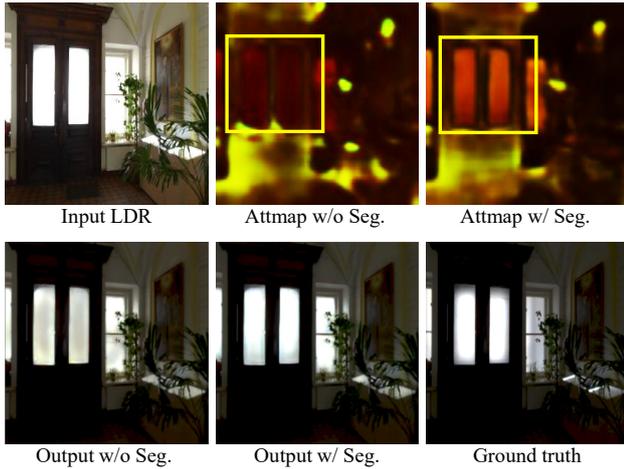
We also provide the visual comparison results in Fig. 8, from which we can see that our proposed LANet recovers the best details and the generated HDR image looks more realistic and more convincing especially viewing under low exposure. To better verify our luminance attention stream and the well-designed LAM, we visualize the attention maps with and without the luminance attention stream as well as the corresponding reconstructed HDR images in Fig. 9. Obviously, with the auxiliary luminance attention stream and LAM, our proposed LANet can recovery better HDR images.

5.4. Comparing LANet with State-of-the-Art

We compare our LANet with one traditional method KOEO [KO14], and five deep learning-based methods, *i.e.*, HDR-CNN [EKD*17], DrTMO [EKM17], Expand-Net [MBRHD18],

Table 3: Quantitative Results on HDR Domain

Method	802 testing pairs				HDR-Eye dataset			
	PU-PSNR \uparrow	PU-SSIM \uparrow	$Q_H\uparrow$	siMSE \downarrow	PU-PSNR \uparrow	PU-SSIM \uparrow	$Q_H\uparrow$	siMSE \downarrow
KOEO [KO14]	26.54	0.870	36.34	2.76	18.91	0.502	29.26	8.95
HDR-CNN [EKD*17]	28.89	0.906	39.97	1.97	20.39	0.572	32.29	6.74
DrTMO [EKM17]	32.12	0.935	39.39	1.55	24.90	0.781	32.34	2.94
Expand-Net [MBRHD18]	27.30	0.862	37.38	2.63	22.29	0.675	30.81	5.02
Santos <i>et al.</i> [SRK20]	27.88	0.886	39.99	2.22	20.35	0.568	31.70	0.652
Liu <i>et al.</i> [LLC*20]	31.21	0.928	39.97	1.43	25.57	0.812	34.62	3.01
LANet	35.79	0.973	42.53	0.56	26.08	0.816	35.63	2.46


Figure 9: An example of attention map generated with and without segmentation learning. Two attention maps are applied a same color boost for visualizing.

Santos *et al.* [SRK20] and Liu *et al.* [LLC*20] on both our self-collected 802 pairs of testing data and the HDR-Eye HDR dataset [NKHE15] for quantitative evaluation. The results are summarized in Table 3, where Q_H is the Q score of HDR-VDP-2 metric, and the results of siMSE metric are in 10^{-1} units. From the table, we can observe that our proposed LANet outperforms the competing methods in terms of Q_H , PU-PSNR, PU-SSIM, and siMSE. To further explain the outperformance of our proposed LANet, we provide some visualization results in Fig. 10. The figure shows that our method can reconstruct a more realistic dynamic range in over-exposed areas and recover the detail of luminance change at non-overexposed areas.

5.5. Application on Image-based Lighting

This section evaluates our panoLANet by applying the reconstructed HDR panoramas to IBL with physically-based rendering. Unlike general images used for display, HDR panoramas mainly provide real environment lighting for physically-based rendering (PBR). Therefore, it is inaccurate to directly evaluate the quality of the panorama by calculating the metrics on images, while we need to evaluate the rendering results rendered by panoramas.

However, for common HDR panoramas, an important issue is that they are not on the same luminance unit. This makes us render

the virtual objects by manually adjusting the light intensity of the panorama. When we perform a quantitative evaluation of a large amount of data, such a manual adjustment will take a long time. What is more, manually adjusting the light intensity is subjective, which will cause the final evaluation results can not be trusted.

Here we can directly use the proposed HDR calibration method to avoid this problem. We calibrate them for each HDR panorama (predicted and ground truth) based on the corresponding LDR panorama in the sIBL dataset. Then we use the calibrated HDR panorama to render a virtual scene under the same rendering configuration. Finally, we make a quantitative comparison of the rendered images to complete the evaluation of our proposed method. Specifically, we use the Blender software with the Cycle rendering engine as the rendering tool for our experiments. The virtual scene we rendered for evaluation consists of four balls with different materials above a diffuse plane, which can be seen in Fig. 11.

Table 4: Comparison of Rendering Results

Method	PSNR \uparrow	SSIM \uparrow	VDP \uparrow	MSE \downarrow
KOEO [KO14]	24.10	0.847	38.55	3.69
HDR-CNN [EKD*17]	23.00	0.831	36.14	4.22
Santos <i>et al.</i> [SRK20]	22.91	0.827	35.26	4.73
Liu <i>et al.</i> [LLC*20]	23.81	0.861	37.77	3.05
U-Net w/ Norm	25.13	0.847	38.48	1.81
U-Net w/ Cal	25.22	0.892	40.43	2.04
U-Net w/ L_{SI}	26.07	0.910	40.68	1.81
panoLANet w/o Att	24.26	0.863	37.84	3.91
LANet	26.20	0.901	41.31	1.71
panoLANet	26.79	0.926	41.66	1.35

Since DrTMO [EKM17] and Expand-Net [MBRHD18] aimed to reconstruct HDR for display, we compare our panoLANet with KOEO [KO14], HDR-CNN [EKD*17], Santos *et al.* [SRK20] and Liu *et al.* [LLC*20] only. In addition, to evaluate the components of our panoLANet, we design a series of variants as follows:

- U-Net w/ Norm: use basic U-Net with the general maximum normalization method.
- U-Net w/ Cal: use basic U-Net with our calibration method.
- U-Net w/ L_{SI} : use basic U-Net with scale-invariant loss.
- LANet: use the same structure as LANet, but applying fine-tune on the panoramas dataset.
- panoLANet w/o Att: use the full panoLANet without skip connection and gated convolution layer.

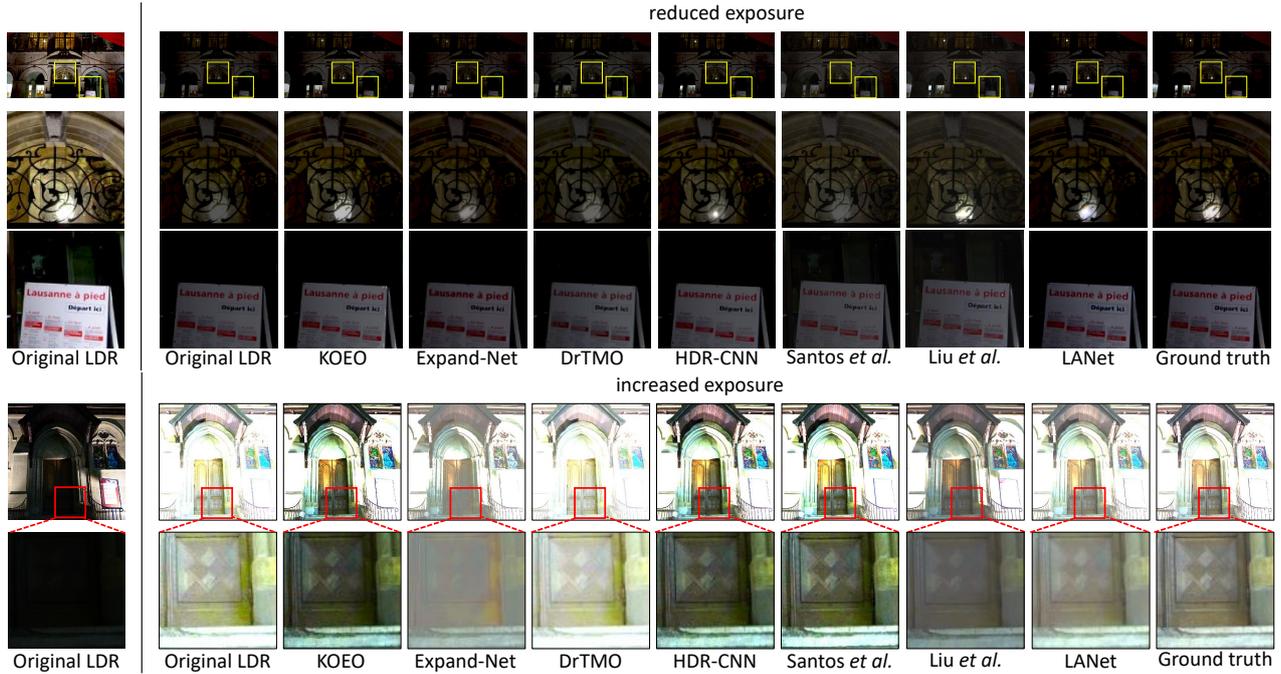


Figure 10: Qualitative results and comparison under different mapping approaches. The comparison on predicted HDR is under the same dynamic range and an approximate exposure for best visual comparison.

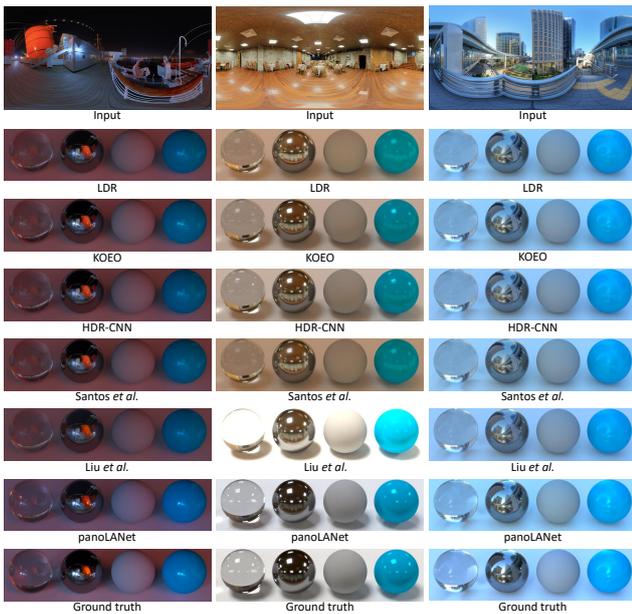


Figure 11: Comparisons on rendering results of predicted panoramas in different scenes. The materials of the four balls from left to right are: glass, glossy, diffuse and mixed diffuse with glossy.

We train the above mentioned variants with panoLANet on the same training data and evaluate their results on the sIBL dataset [Hdr12]. The results are summarized in Table 4 and the results of siMSE metric are in 10^{-3} units. From the results, we can

see that our panoLANet gets the best results in all the metrics, especially the result of SSIM is as high as 0.926. In addition, the HDR calibration algorithm is also superior to the general maximum normalization method in the HDR reconstruction of panorama, which again indicates the advantages of our calibration method.

In order to demonstrate the advantages of panoLANet more clearly, we show the comparison results of our method with the previous methods in three different scenes: night scene, indoor and outdoor scenes. As can be seen from the qualitative comparison results in Fig. 11, our method estimates the position and brightness of the light source more accurately, which is reflected in the more realistic highlight position and brightness of the render results. In addition, the \mathcal{L}_{SI} is adopted to promote a more accurate ratio relationship between the intensity of the light source and the background, which can be observed from the results in the second column of Fig. 11.

5.6. Discussion

Although our method can reconstruct the high dynamic range well in many scenes, it still fails in some cases, especially in extremely over-exposure scenes. As shown in Fig. 12, when there is a large over-exposure region in the image, our method cannot restore texture information and luminance information correctly.

Regarding runtime, note that our training stage is carried out off-line, and here we only offer the runtime for the testing stage under our experimental environment. For per LDR image of the size 1920×1080 and the panorama of the size 1024×512 , it requires about 3 seconds and 1 second respectively to complete the testing throughout the trained model.

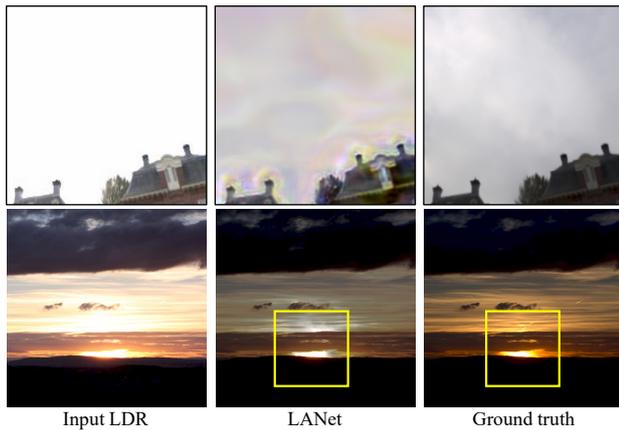


Figure 12: Failure cases. The top three images are with highly over-exposed areas that the textual detail cannot be recovered well. The bottom three images are without color information around the overexposed area, which leads to color information missing in the input LDR image for inferring HDR image.

6. Conclusion

In this paper, we propose an end-to-end and trainable luminance attentive network for HDR image reconstruction from single LDR images, as well as its extended network for HDR panorama reconstruction for image-based lighting usage. Rather than using the general maximum normalization method on HDR data, we calibrate original HDR images to the similar luminance scale corresponding to LDR images. This treatment gives us many benefits on HDR reconstruction tasks, such as obtaining a set of HDR images with a similar luminance scale and getting pixel-level luminance segmentation labels automatically without requiring any extra manual annotation. Also, the designed luminance attention module can well explore the estimated luminance semantic segmentation to pay attention to over-exposure and under-exposure areas for better reconstructing the HDR images. The extended panoLANet also achieves a better performance on HDR panorama reconstruction for IBL usage. Our future work includes combining deep learning-based HDR reconstruction with image inpainting to perform image restoration in the HDR domain and solve the texture loss of LDR images in over-exposed areas with no light source.

Acknowledgement

This work was partly supported by the Key Technological Innovation Projects of Hubei Province (2018AAA062), NSFC (NO. 61972298, 61872277, 62171324), and CAAI-Huawei MindSpore Open Fund.

References

[AFR*07] AKYUZ A. O., FLEMING R., RIECKE B. E., REINHARD E., BULTHOFF H. H.: Do hdr displays support ldr content?: a psychophysical evaluation. *ACM Trans. Graph.* 26, 3 (2007), 38–49. 1, 2

[AMS08] AYDIN T. O., MANTIUK R., SEIDEL H.-P.: Extending quality metrics to full luminance range images. In *P SPIE IST ELECT IM* (2008), p. 68060B. 7

[BLDC06] BANTERLE F., LEDDA P., DEBATTISTA K., CHALMERS A.: Inverse tone mapping. In *GRAPHITE* (2006), pp. 349–356. 2

[CLZX21] CHEN Z., LONG C., ZHANG L., XIAO C.: Canet: A context-aware network for shadow removal. In *ICCV* (2021). 2

[CMY19] CHEN C., MCCLOSKEY S., YU J.: Analyzing modern camera response functions. In *WACV* (2019), pp. 1961–1969. 7

[DLMM16] DUFAUX F., LE CALLET P., MANTIUK R. K., MRAK M.: Chapter 7 - evaluation of tone mapping operators for hdr video. In *High Dynamic Range Video*. 2016, pp. 185–207. 6

[DLXX21] DONG X., LONG C., XU W., XIAO C.: Dual graph convolutional networks with transformer and curriculum learning for image captioning. In *ACM Multimedia* (2021). 2

[DLZX19] DING B., LONG C., ZHANG L., XIAO C.: Argan: Attentive recurrent generative adversarial network for shadow detection and removal. In *ICCV* (2019), pp. 10213–10222. 2

[DXO19] DXOMARK: Cameras sensor database. <https://www.dxomark.com/Cameras/>, 2019. 7

[EKD*17] EILERTSEN G., KRONANDER J., DENES G., MANTIUK R. K., UNGER J.: Hdr image reconstruction from a single exposure using deep cnns. *ACM Trans. Graph.* 36, 6 (2017), 178–188. 1, 3, 4, 7, 8, 9

[EKM17] ENDO Y., KANAMORI Y., MITANI J.: Deep reverse tone mapping. *ACM Trans. Graph.* 36, 6 (2017), 177–188. 1, 3, 8, 9

[EPF14] EIGEN D., PUHRSCH C., FERGUS R.: Depth map prediction from a single image using a multi-scale deep network. In *NIPS* (2014), pp. 2366–2374. 2, 7

[Fai17] FAIRCHILD M. D.: The hdr photographic survey. <http://rit-mcsl.org/fairchild//HDR.html>, 2017. 6

[FS00] FUNT B., SHI L.: Funt et al. hdr dataset. https://www2.cs.sfu.ca/~colour/data/funt_hdr/, 2000. 6

[GJAF09] GROSSE R., JOHNSON M. K., ADELSON E. H., FREEMAN W. T.: Ground truth dataset and baseline evaluations for intrinsic image algorithms. In *ICCV* (2009), pp. 2335–2342. 4

[GSY*17] GARDNER M.-A., SUNKAVALLI K., YUMER E., SHEN X., GAMBARETTO E., GAGNÉ C., LALONDE J.-F.: Learning to predict indoor illumination from a single image. *ACM Trans. Graph.* 36, 6 (2017), 176. 6, 7

[Hdr12] HDRLABS: sibl archive: Free hdri sets for smart image-based lighting. <http://www.hdrilabs.com/sibl/archive.html>, 2012. 2, 6, 7, 10

[HGAL19] HOLD-GEOFFROY Y., ATHAWALE A., LALONDE J.-F.: Deep sky modeling for single image outdoor lighting estimation. In *CVPR* (2019), pp. 6927–6935. 6, 7

[HKU15] HAJISHARIF S., KRONANDER J., UNGER J.: Adaptive dualiso HDR reconstruction. *EURASIP J. Image Video Process.* 2015 (2015), 41. 2

[HLX21a] HU T., LONG C., XIAO C.: Crd-cgan: Category-consistent and relativistic constraints for diverse text-to-image generation. *arXiv preprint arXiv:2107.13516* (2021). 2

[HLX21b] HU T., LONG C., XIAO C.: A novel visual representation on text using diverse conditional gan for visual recognition. *TIP* 30 (2021), 3499–3512. 2

[HLYG18] HUA G., LONG C., YANG M., GAO Y.: Collaborative active visual recognition from crowds: A distributed ensemble approach. *T-PAMI* 40, 3 (2018), 582–594. 2

[HZR16] HE K., ZHANG X., REN S.: Identity mappings in deep residual networks. In *ECCV* (2016), pp. 630–645. 4

[HZRS16] HE K., ZHANG X., REN S., SUN J.: Deep residual learning for image recognition. In *CVPR* (2016), pp. 770–778. 4

[ILBH20] ISLAM A., LONG C., BASHARAT A., HOOGS A.: Doagan: Dual-order attentive generative adversarial network for image copy-move forgery detection and localization. In *CVPR* (2020). 2

- [ILR21] ISLAM A., LONG C., RADKE R.: A hybrid attention mechanism for weakly-supervised temporal action localization. In *AAAI* (2021). 3
- [IS15] IOFFE S., SZEGEDY C.: Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML* (2015), pp. 448–456. 4
- [KB15] KINGMA D. P., BA J.: Adam: A method for stochastic optimization. In *ICLR* (2015). 7
- [KGB*14] KRONANDER J., GUSTAVSON S., BONNET G., YNNERMAN A., UNGER J.: A unified framework for multi-sensor hdr video reconstruction. *Signal Process. Image Commun.* 29, 2 (2014), 203–215. 6
- [KGBU13] KRONANDER J., GUSTAVSON S., BONNET G., UNGER J.: Unified HDR reconstruction from raw CFA data. In *ICCP* (2013), pp. 1–9. 6
- [KO14] KOVALESKI R. P., OLIVEIRA M. M.: High-quality reverse tone mapping for a wide range of exposures. In *SIBGRAP* (2014), pp. 49–56. 1, 2, 8, 9
- [KOK19] KIM S. Y., OH J., KIM M.: Deep sr-itm: Joint learning of super-resolution and inverse tone-mapping for 4k uhd hdr applications. In *ICCV* (2019), pp. 3116–3125. 3
- [KTC12] KUO P.-H., TANG C.-S., CHIEN S.-Y.: Content-adaptive inverse tone mapping. In *VCIP* (2012), pp. 1–6. 1
- [LAK18] LEE S., AN G. H., KANG S.-J.: Deep chain hdri: Reconstructing a high dynamic range image from a single low dynamic range image. *IEEE Access* 6 (2018), 49913–49924. 3
- [LH15] LONG C., HUA G.: Multi-class multi-annotator active learning with robust gaussian process for visual recognition. In *ICCV* (2015). 2
- [LH17] LONG C., HUA G.: Correlational gaussian processes for cross-domain visual recognition. In *CVPR* (2017). 2
- [LHAK18] LEE S., HWAN AN G., KANG S.-J.: Deep recursive hdri: Inverse tone mapping using generative adversarial networks. In *ECCV* (2018), pp. 596–611. 1, 3
- [LLC*20] LIU Y.-L., LAI W.-S., CHEN Y.-S., KAO Y.-L., YANG M.-H., CHUANG Y.-Y., HUANG J.-B.: Single-image hdr reconstruction by learning to reverse the camera pipeline. In *CVPR* (2020), pp. 1651–1660. 3, 9
- [LLZ*20] LIU D., LONG C., ZHANG H., YU H., DONG X., XIAO C.: Arshadowgan: Shadow generative adversarial network for augmented reality in single light scenes. In *CVPR* (2020), pp. 8136–8145. 3
- [LWH*14] LONG C., WANG X., HUA G., YANG M., LIN Y.: Accurate object detection with location relaxation and regionlets re-localization. In *ACCV* (2014). 2
- [MAF*09] MASIA B., AGUSTIN S., FLEMING R. W., SORKINE O., GUTIERREZ D.: Evaluation of reverse tone mapping through varying exposure conditions. *ACM Trans. Graph.* 28, 5 (2009), 160–172. 1, 2
- [MBRHD18] MARNERIDES D., BASHFORD-ROGERS T., HATCHETT J., DEBATTISTA K.: Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. In *CGF* (2018), vol. 37(2), pp. 37–49. 1, 3, 8, 9
- [MDK08] MANTIUK R., DALY S., KEROFKY L.: Display adaptive tone mapping. *ACM Trans. Graph.* 27, 3 (2008), 68–79. 7
- [MMP*07] MCGUIRE M., MATUSIK W., PFISTER H., CHEN B., HUGHES J. F., NAYAR S. K.: Optical splitting trees for high-precision monocular imaging. *Computer Graphics and Applications* 27, 2 (2007), 32–42. 2
- [MSG17] MASIA B., SERRANO A., GUTIERREZ D.: Dynamic range expansion based on image statistics. *Multim. Tools Appl.* 76, 1 (2017), 631–648. 1
- [NKHE15] NEMOTO H., KORSHUNOV P., HANHART P., EBRAHIMI T.: Visual attention in ldr and hdr images. In *VPQM* (2015). 2, 6, 7, 9
- [NMDSL15] NARWARIA M., MANTIUK R., DA SILVA M. P., LE CALLET P.: Hdr-vdp-2.2: a calibrated method for objective quality prediction of high-dynamic range and standard images. *Electronic Imaging* 24, 1 (2015), 501–514. 7
- [ODO16] ODENA A., DUMOULIN V., OLAH C.: Deconvolution and checkerboard artifacts. *Distill* (2016). 4
- [Pul14] PULLI K.: Flexisp: a flexible camera image processing framework. *ACM Trans. Graph.* 33, 6 (2014), 231:1–231:13. 2
- [RFB15] RONNEBERGER O., FISCHER P., BROX T.: U-net: Convolutional networks for biomedical image segmentation. In *MICCAI* (2015), pp. 234–241. 4
- [RTS*07] REMPEL A. G., TRENTACOSTE M., SEETZEN H., YOUNG H. D., HEIDRICH W., WHITEHEAD L., WARD G.: Ldr2hdr: on-the-fly reverse tone mapping of legacy video and photographs. *ACM Trans. Graph.* 26, 3 (2007), 39–50. 1, 2
- [SRK20] SANTOS M. S., REN T. I., KALANTARI N. K.: Single image HDR reconstruction using a CNN with masked features and perceptual loss. *ACM Trans. Graph.* 39, 4 (2020), 80. 9
- [SWL*21] SHI L., WANG L., LONG C., ZHOU S., ZHOU M., NIU Z., HUA G.: Sgc: Sparse graph convolution for pedestrian trajectory prediction. In *CVPR* (2021). 3
- [TKTS11] TOCCI M. D., KISER C., TOCCI N., SEN P.: A versatile HDR video production system. *ACM Trans. Graph.* 30, 4 (2011), 41. 2
- [UUVL16] ULYANOV D., VEDALDI A., LEMPITSKY V.: Instance normalization: The missing ingredient for fast stylization. *CoRR abs/1607.08022* (2016). 4
- [War19] WARD G.: High dynamic range image examples. <http://www.anyhere.com/gward/hdrenc/pages/originals.html>, 2019. 6
- [WLZX19] WEI J., LONG C., ZOU H., XIAO C.: Shadow inpainting and removal using generative adversarial networks with slice convolutions. *CGF* 38, 7 (2019), 381–392. 2
- [WZW19] WANG C., ZHAO Y., WANG R.: Deep inverse tone mapping for compressed images. *IEEE Access* 7 (2019), 74558–74569. 1
- [XDCW02] XIAO F., DICARLO J. M., CATRYSSSE P. B., WANDELL B. A.: High dynamic range imaging of natural scenes. In *CIC* (2002), pp. 337–342. 6
- [YLY*19] YU J., LIN Z., YANG J., SHEN X., LU X., HUANG T. S.: Free-form image inpainting with gated convolution. In *ICCV* (2019), pp. 4471–4480. 6
- [ZA21] ZHANG Y., AYDIN T. O.: Deep HDR estimation with generative detail reconstruction. *Comput. Graph. Forum* 40, 2 (2021), 179–190. 3
- [Zaa19] ZAAL G.: Hdri haven. <https://hdrihaven.com/>, 2019. 6
- [ZL17] ZHANG J., LALONDE J.-F.: Learning high dynamic range from outdoor panoramas. In *ICCV* (2017), pp. 4519–4528. 3
- [ZLW19] ZHANG S., LIANG R., WANG M.: Shadowgan: Shadow synthesis for virtual objects with conditional adversarial networks. *Comput. Vis. Media* 5, 1 (2019), 105–115. 3
- [ZLW*20] ZHANG J., LONG C., WANG Y., YANG X., MEI H., YIN B.: Multi-context and enhanced reconstruction network for single image super resolution. In *ICME* (2020). 2
- [ZLW*21] ZHANG J., LONG C., WANG Y., PIAO H., MEI H., YANG X., YIN B. Y.: A two-stage attentive network for single image super resolution. *TCSVT* (2021). 2
- [ZLY*20] ZHANG L., LONG C., YAN Q., ZHANG X., XIAO C.: Clagan: A context and lightness aware generative adversarial network for shadow removal. *CGF* 39, 7 (2020). 2
- [ZLZX20] ZHANG L., LONG C., ZHANG X., XIAO C.: Ris-gan: Explore residual and illumination with generative adversarial networks for shadow removal. In *AAAI* (2020). 2