

Joint Multi-Scale Tone Mapping and Denoising for HDR Image Enhancement

Litao Hu
Purdue University
hu430@purdue.edu

Huaijin Chen
Sensebrain Technology
chenhuaijin@sensebrain.ai

Jan P. Allebach
Purdue University
allebach@purdue.edu

Abstract

An image processing unit (IPU), or image signal processor (ISP) for high dynamic range (HDR) imaging usually consists of demosaicing, white balancing, lens shading correction, color correction, denoising, and tone-mapping. Besides noise from the imaging sensors, almost every step in the ISP introduces or amplifies noise in different ways, and denoising operators are designed to reduce the noise from these sources. Designed for dynamic range compressing, tone-mapping operators in an ISP can significantly amplify the noise level, especially for images captured in low-light conditions, making denoising very difficult. Therefore, we propose a joint multi-scale denoising and tone-mapping framework that is designed with both operations in mind for HDR images. Our joint network is trained in an end-to-end format that optimizes both operators together, to prevent the tone-mapping operator from overwhelming the denoising operator. Our model outperforms existing HDR denoising and tone-mapping operators both quantitatively and qualitatively on most of our benchmarking datasets.

1. Introduction

HDR image enhancement is a complex task consisting of two major sub-tasks, namely image tone mapping and image denoising. In a typical ISP for HDR images, tone-mapping operators (TMOs) and image denoising operators are separate components. A tone-mapping operator aims to compress the dynamic range of the HDR image by brightening the low-lights and dimming the high-lights so that the HDR content can be viewed on low-dynamic-range displays preserving its perceptual contrast. Many HDR images are captured under low-exposure settings or low-light conditions, to capture and preserve a wider dynamic range without over-exposing the highlight regions. As a result, noise in HDR images can be a very challenging problem for any existing denoising algorithms designed for standard-dynamic-range (SDR) images. Meanwhile, denoising often requires the inputs to be in a linear color space, which can hardly hold true due to the multi-exposure HDR capture and

non-linear tone-mapping.

With the denoising and tone-mapping treated as separate modules in conventional ISPs, it is natural for us to wonder which operator should be applied first. Denoising after TMOs [15] can pose significant challenges to denoising, since TMOs are usually highly non-linear and tends to amplify the noise. Denoising before TMOs, in contrast, may cause the remaining noise after denoising being significantly amplified by the subsequent TMOs, requiring extra denoising steps. Therefore, it makes sense for us to design the denoising and tone-mapping jointly. Meanwhile, we observe that both denoising and TMOs can benefit from multi-scale processing. For instance, tone-mapping can be achieved by adjusting the brightness to the high- and low-frequency component of the HDR image differently [10], and denoising on decomposed multi-scale inputs also receives better results than the single-scale counterpart [36].

Thus, in this paper, we propose to *jointly* optimize the two operators within a *multi-scale* network for HDR image enhancement and showcase that our proposed framework achieves better results than state-of-the-art HDR tone-mapping methods and separately performed denoising and tone-mapping procedures on recent HDR benchmark datasets both quantitatively and qualitatively. Our key contributions are the following:

- We propose a novel learned differentiable single image denoising module based on discrete cosine transform (DCT);
- We propose a multi-scale image enhancement framework that jointly optimizes the TMOs and denoising operators for HDR image enhancement;
- We investigate the effect of different orderings of the TMOs and denoising operators and demonstrate that having TMOs applied first yields better final results.

2. Related Work

2.1. HDR Image Tone Mapping

Tone-mapping algorithms can be categorized into learning-based methods and traditional methods depend-

ing on whether training is needed. With the advancement in deep learning, learning-based TMOs are become more advantageous, especially in challenging cases. Within the learning-based category, methods can be further divided into supervised [7, 44, 27, 42, 6, 21, 54, 37, 43, 46, 12, 56, 39, 33, 28], semi-supervised [47], and unsupervised [52, 22, 16, 55, 41, 32], depending on whether ground-truth is provided during training. In general, most supervised TMOs outperform semi-supervised and unsupervised TMOs, while the latter requires fewer paired training samples and thus can be easily implemented in a wide range of applications.

In particular, multi-scale structures such as the Laplacian pyramid has been used by many learning-based image enhancement algorithms including [9], [24], and [1].

In our work, we adopt the CSRNet [20] as our backbone tone-mapping network for its decent performance, lightweight, and fully convolutional structure that makes it compatible with any size input.

2.2. HDR Image Denoising

Deep learning has been applied extensively to the image denoising field [38], outperforming many of the traditional denoising algorithms. Based on the number of images used as input for a given scene, these deep learning-based image denoising methods can be categorized into single-image denoising and multi-image denoising. Single-image denoising [50, 51, 17, 8, 25, 30, 3, 31, 40, 2] aims to denoise a target scene with only a single shot from the scene. On the other hand, multi-image denoising [14] aims to denoise a target scene by utilizing multiple shots from the same scene. A common strategy in multi-image denoising utilizes images from a shot burst, providing a richer spatial-temporal context for the model. However, this requires image registration preprocessors for aligning and merging multiple images from a shot burst, which can cause other artifacts and is usually computationally intensive.

Based on whether a multi-scale structure is used, denoising methods can also be classified into single-scale methods and multi-scale methods. Multi-scale denoising has been an effective strategy for extracting and providing multi-scale information to the denoising models [30, 3, 40].

In our work, we adopt a single-image multi-scale denoising framework and propose a novel deep DCT denoising module inspired by the traditional DCT denoising algorithm.

2.3. Joint Tone Mapping and Denoising for HDR Images

For a noisy raw HDR image, different modules in an ISP will have different effects on the noise distribution. Modules such as demosaicing and tone-mapping can sometimes boost the noise significantly if not designed wisely. There-

fore, some recent works try to combine denoising with these modules, to get a better measurement of the noise distribution and provide more intermediate information to the denoising algorithm. [13, 23, 11, 29] proposed to jointly perform image demosaicing and image denoising. However, a major drawback of these methods is that since demosaicing is usually performed at the beginning of an ISP, the remaining noise after the joint denoising can still be amplified by the following operations, especially those highly non-linear ones such as tone mapping. As a result, extra denoising modules are often needed to clean up the noise remaining from these operations. A joint tone mapping and denoising framework, on the other hand, does not have such a drawback since TMOs are usually placed near the end of an ISP pipeline. Moreover, by jointly optimizing two modules, the TMO can provide information and context on where the noise gets amplified more. [15] has proved the effectiveness of applying joint tone mapping and denoising on HDR contents.

In our work, we extend the idea of joint tone mapping and denoising to a multi-scale framework, where the TMOs and denoising operators are jointly optimized at multiple scales in a Laplacian pyramid.

3. Proposed Method

In this section, we will present our joint multi-scale tone mapping and denoising framework for an HDR image in detail. The overall structure of our proposed framework is shown in Figure 1.

3.1. Deep DCT Denoising for a Single HDR Image

We will first introduce our deep DCT denoising network designed for single HDR image denoising. The architecture of our deep DCT denoising algorithm for a single HDR image is shown in Figure 1.

DCT denoising was first proposed in 2011 by [49] and was widely used because of its simplicity and effectiveness. Later in 2017, [36] extended the DCT denoising algorithm to a multi-scale setting. The deep DCT denoising network we propose is a learnable DCT denoising algorithm that combines the simplicity of the traditional DCT denoising algorithm and the effectiveness of modern deep neural networks.

In traditional DCT denoising algorithms, a preset threshold is chosen to zero out DCT coefficients below the threshold in order to create sparsity in the DCT domain of local image tiles. However, this operation is not differentiable and thus is not suitable for our learning-based framework. Therefore, instead of applying a learned threshold to the DCT coefficients, we multiply the DCT coefficients with learned re-weighting maps with re-weighting values between 0 and 1. In a re-weighting map, a value close to 0

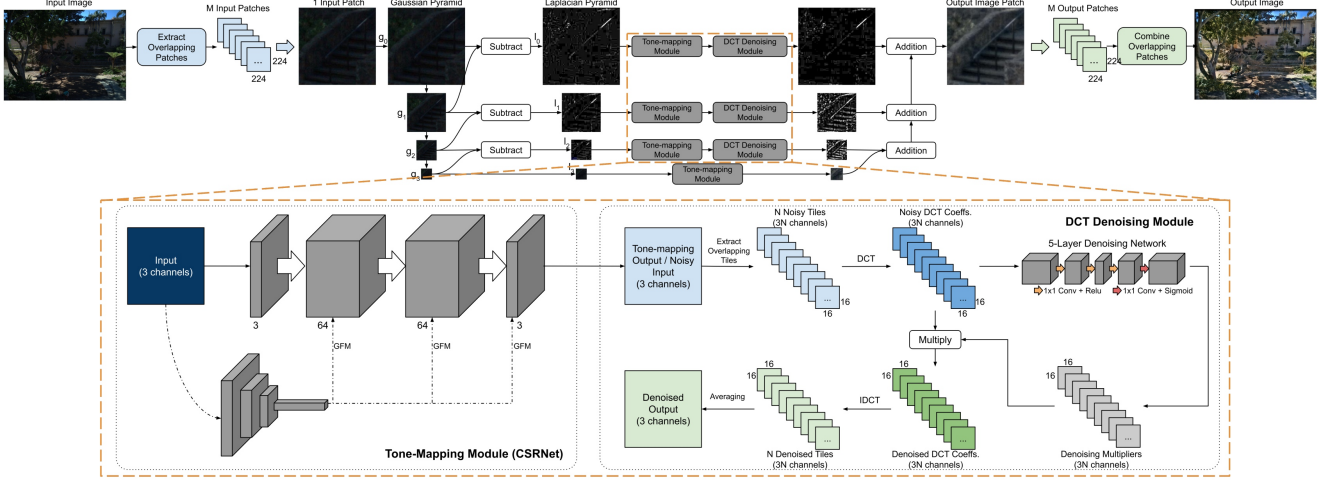


Figure 1: Our proposed HDR image enhancement framework. The upper portion of the figure shows a high-level workflow of our framework and the multi-scale structure. The lower portion of the figure shows more details of the tone-mapping module and the denoising module.

is equivalent to the case where the corresponding DCT coefficient is below the threshold, while a value close to 1 is equivalent to the case where the DCT coefficient is above the threshold. Using re-weighting maps not only allows for differentiable parameters but also provides more flexibility to the denoising module, leading to better denoising performance.

The input to our DCT denoising module is a 3-channel noisy image patch in RGB. As the first step in the denoising module, the input image patch will be split into overlapping tiles in size of 16×16 , which will then be reshaped into a stack of tiles. We adopt 16×16 as the size of all local image tiles following [49]. Organizing the tiles this way allows us to perform the DCT easily for every tile. These noisy DCT tiles will then be passed to a 5-layer CNN to estimate a denoising multiplier within the DCT domain, which will then be multiplied element-by-element by the noisy DCT tiles to get the same number of denoised DCT tiles. These denoised DCT tiles will then go through an inverse discrete cosine transform (IDCT), as well as a tile-averaging operation, to reconstruct the denoised output image patch. Note that the tile-averaging operation here is a simple averaging operation applied to the overlapping pixels among neighboring tiles. For example, if a pixel in the output image patch is overlapped by 4 neighboring tiles that reconstruct this output image patch, the value for this pixel will then be the average of the 4 values of the corresponding locations in these 4 overlapping tiles.

Our DCT denoising process can be represented by Equation 1, where C_{noisy} and $C_{denoised}$ represent the DCT coefficients of the noisy input and the denoised output, respectively. \mathcal{F} represents the non-linear function learned by the

5-layer CNN for estimating the denoising multipliers, and \odot represents element-by-element multiplication. Note that the 5-layer CNN here for estimating the denoising multipliers contains only 1×1 2D convolutional layers. We made this design choice intentionally because we believe that in the DCT domain, the coefficients at different frequencies should be considered separately, while coefficients at the same frequency should be grouped together. Thus, convolutional kernels that are commonly used in the image domain, such as a 3×3 2D convolutional layer, would not be suitable here.

$$C_{denoised} = C_{noisy} \odot \mathcal{F}(C_{noisy}) \quad (1)$$

3.2. Joint Multi-Scale Tone-mapping and Denoising Framework

The other important component of our proposed framework is its joint multi-scale tone-mapping and denoising framework. The overall structure of our framework and a closer look at the multi-scale structure are shown in Figure 1.

Multi-scale structures have been used extensively in many image processing and computer vision tasks for their effectiveness in producing finer results than their single-scale counterparts. Some [3, 40] incorporate the multi-scale aspect by simply downsampling the input a few times, and then processing each scale separately, while others [30, 36] utilize domain transformations such as the discrete wavelet transform (DWT) and DCT.

In our proposed framework, we utilize the Laplacian pyramid to obtain multiple scales of an input image. To be specific, a noisy input RGB image will first be divided into

half-overlapping image patches of size 224×224 pixels. Then every image patch will be decomposed into four Gaussian pyramid layers (g_0, g_1, g_2, g_3). From there, three Laplacian layers (l_0, l_1, l_2) will be computed from the neighboring Gaussian layers, which then forms a four-layer Laplacian pyramid along with the Gaussian base layer (g_3 or l_3).

From the Laplacian pyramid, the Laplacian layers l_0, l_1 , and l_2 will be passed to their corresponding tone mapping networks, followed by their corresponding deep DCT denoising modules, while the base layer (g_3 or l_3) will only be tone-mapped. We don't denoise the base layer because most high-frequency noise will be separated to the upper layers and it is unnecessary for a denoising operator in the base layer. This can help reduce the model size and speed up inference.

These tone-mapped and denoised output Laplacian layers will then be combined to reconstruct the corresponding tone-mapped and denoised output RGB image patch, following the inverse process of obtaining the Laplacian pyramid. Finally, all the output image tiles will be combined into a single output RGB image using a modified raised-cosine filter, following the same way of merging overlapping image tiles in [35]. As shown in our results, this modified raised-cosine filter does a very good job merging image tiles without showing any obvious artifacts.

3.3. Training Scheme and Loss Functions

To better optimize both the tone-mapping networks and the denoising modules, we adopt a three-phase training scheme in our experiments. To be specific, we train the model in three consecutive phases, namely tone-mapping training, denoising training, and joint training.

In each training phase, a different part of the joint framework will be trained and different loss functions will be used to serve each training purpose.

The first training phase, i.e. the tone-mapping training phase, focuses on pre-training the tone-mapping networks. In this phase, the denoising modules will be temporarily moved out of the joint framework, meaning that each layer in the pyramid will only go through a tone-mapping network. The loss function for this phase is computed according to Equations 2 and 3, where $L_{l_0}, L_{l_1}, L_{l_2}, L_{l_3}$ are the L1 loss computed at each Laplacian layer, and $l_i^{(1)}$ and l_i represent, respectively, the tone-mapped Laplacian layer and the corresponding Laplacian layer from the tone-mapped ground-truth image. Here, $\lambda_0, \lambda_1, \lambda_2$, and λ_3 are hyper-parameters, and the selected values for them can be found in Section 4.1.

$$L_{p1} = \lambda_0 L_{l_0}^{(1)} + \lambda_1 L_{l_1}^{(1)} + \lambda_2 L_{l_2}^{(1)} + \lambda_3 L_{l_3}^{(1)} \quad (2)$$

$$L_{l_i}^{(1)} = |l_i^{(1)} - l_i|, \quad i = 0, 1, 2, 3 \quad (3)$$

The second phase, i.e. the denoising phase, focuses on

pre-training the denoising networks. The denoising modules are added back to the joint framework during this phase, while the pre-trained tone-mapping networks are fixed. The loss function for this phase is computed according to Equations 4-6, where $L_{l_0}, L_{l_1}, L_{l_2}$ are the L1 loss computed at each Laplacian except for l_3 , and $l_i^{(0)}$ and $l_i^{(2)}$ represent, respectively, the input Laplacian layer and the tone-mapped denoised Laplacian layer. $L_{denoise}$ measures the denoising loss by computing the L1 loss between the reconstructed denoised output and the clean ground-truth image without tone-mapping (Y_{merged}). Here, we use R to represent the reconstruction operator. Again, $\lambda_0, \lambda_1, \lambda_2$, and λ_d are hyper-parameters, and the selected values for them can be found in Section 4.1.

$$L_{p2} = \lambda_0 L_{l_0}^{(2)} + \lambda_1 L_{l_1}^{(2)} + \lambda_2 L_{l_2}^{(2)} + \lambda_d L_{denoise} \quad (4)$$

$$L_{l_i}^{(2)} = |l_i^{(2)} - l_i|, \quad i = 0, 1, 2 \quad (5)$$

$$L_{denoise} = |R(l_0^{(2)}, l_1^{(2)}, l_2^{(2)}, l_3^{(0)}) - Y_{merged}| \quad (6)$$

The last phase, i.e. the joint training phase, trains both the tone-mapping networks and the denoising modules together. The loss function for this phase is simply an L1 loss computed between the reconstructed output from the tone-mapped denoised Laplacian layers and the clean tone-mapped ground-truth image Y_{final} , as shown in Equation 7.

$$L_{p3} = |R(l_0^{(2)}, l_1^{(2)}, l_2^{(2)}, l_3^{(1)}) - Y_{final}| \quad (7)$$

4. Experiments

In this section, we will demonstrate the effectiveness of our proposed framework through both qualitative and quantitative evaluation on several HDR image datasets.

4.1. Training Settings and Dataset

Our experiments are run on a single Nvidia Tesla GPU. Detailed training settings for the three training phases for our joint framework are summarized in Table 1. Adam optimizers with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ are adopted for all 3 training phases. We empirically choose the hyper-parameters in the loss functions (Equation 2 and Equation 4) to be $\lambda_0 = 2, \lambda_1 = 2, \lambda_2 = 2, \lambda_3 = 1, \lambda_d = 1$.

Phase	Batch Size	Learning Rate	Epochs
1	16	1e-5	500
2	16	1e-5	500
3	16	1e-6	1000

Table 1: Training settings for the three phases.




Figure 2: Qualitative comparisons between methods on examples from the HDR+ testing set. Four sub-regions selected from each example are zoomed in 10 times to show more details. Images here have been down-sampled for better viewing experience.

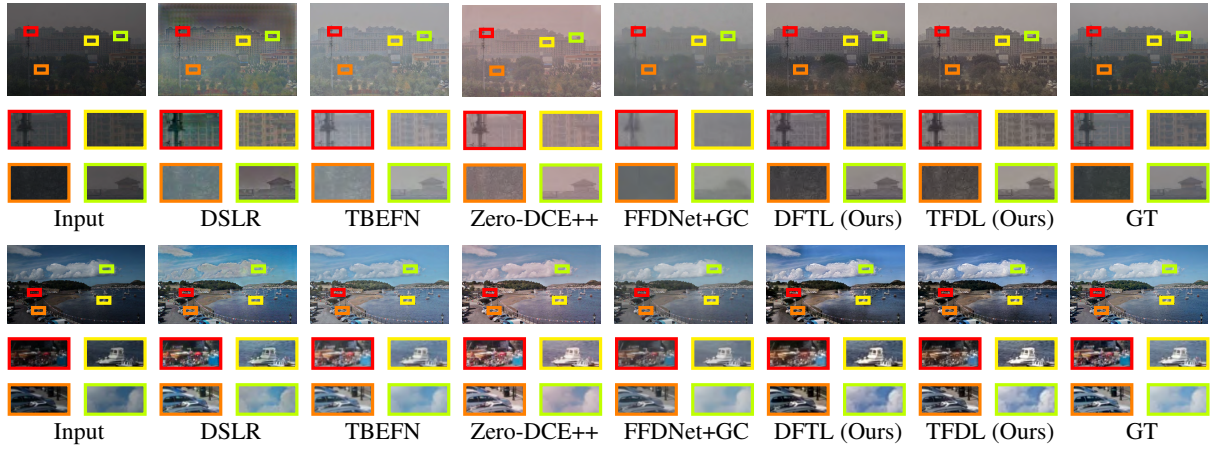


Figure 3: Qualitative comparisons between methods on examples from the LOL dataset (the first row) and SCIE dataset (the second row). Four sub-regions selected from each example are zoomed in 5 times to show more details.

We use the Google HDR+ dataset [19] as our training dataset, as it provides both the merged intermediate ground-truth image (denoised using burst captures) and the final ground-truth image (denoised and tone-mapped) for calculating our loss functions. From the HDR+ dataset, we randomly select 2000 scenes for training, 200 scenes for validation, and 1300 scenes for testing and benchmarking.

The input images we used from the HDR+ dataset are the reference frames from the N input burst photos of each individual scene, and the merged intermediate ground-truth images are the intermediate raw results of aligning and merging their corresponding input burst. For simplicity, we pre-processed the raw input images and raw merged interme-

diate images by demosaicing, color correction, and gamma correction to convert them to sRGB color space before using them. Note that the ground-truth images in the HDR+ dataset are obtained through a series of processing steps in the finishing pipelines besides denoising and tone-mapping. However, including all the finishing steps in our framework would overly complicate our method, so the rest of the processing steps in the finishing pipeline are not explicitly performed in our framework when obtaining the final tone-mapped and denoised output, and we can assume that these operations are implicitly learned through the tone-mapping modules and denoising modules in our joint multi-scale framework.

Since the provided final ground-truth images are cropped, we need to align the input images and the merged intermediate ground-truth images with their final counterparts. For convenience, we used the Scale-Invariant Feature Transform (SIFT) [4] to align and crop the input and intermediate images as a preprocessing step.

Our joint framework takes input with a size of 224×224 pixels, so we cannot use the full images to train our model. Instead, we split all 2000 input images into tiles of size 224×224 pixels, along with their intermediate merged ground-truth image and final ground-truth image.

4.2. Results

4.2.1 Full-Reference Comparison

To quantitatively evaluate the performance of our proposed joint tone-mapping and denoising framework, we measured the SSIM, PSNR, LPIPS [53], and Tone-Mapping Quality Index (TMQI) [48] on the HDR+ testing set (1300 scenes), as well as the LOL training dataset (485 scenes) [45] and part 1 of the SCIE dataset (360 scenes) [5]. Note that our models are neither trained nor fine-tuned on any subsets of the LOL dataset or the SCIE dataset. For the HDR+ dataset, we used the reference frame as the input and for the SCIE dataset, we used the first image (the image with the lowest exposure) of every scene as the input. We tested several state-of-the-art image enhancement methods including DSLR[28], TBEFN[32], and Zero-DCE++[26] as a comparison to our method, and we also compared the denoising performance of our method to one of the state-of-the-art denoising methods named FFDNet[51]. Note that among these methods, DSLR, TBEFN, and FFDNet are supervised methods, while Zero-DCE++, an extension to its original version Zero-DCE[16], is trained through unsupervised learning. Note that for DSLR, TBEFN, Zero-DCE++, and FFDNet, we used their corresponding pre-trained versions in all our experiments without fine-tuning. For FFDNet, since it is merely a denoising algorithm, we paired it with a standard gamma correction (GC) with γ set to 2.2.

To see whether performing tone-mapping or denoising first would produce better results, we also trained our models in both configurations. For convenience, we use “DFTL” (denoise first, tone-map last) to represent the model where tone-mapping is performed after denoising, and “TFDL” (tone-map first, denoise last) to represent the model where tone-mapping is performed before denoising. The results are summarized in Table 2, Table 3, and Table 4. From these quantitative results, we can see that our method outperforms the other methods on most of the benchmarking datasets, and is significantly advantageous on the HDR+ testing set. This is likely due to the fact that, among these datasets we tested, images from the HDR+ testing set contain heavy noise, while images from the LOL dataset and SCIE dataset contain only minimal noise.

Methods	SSIM	PSNR	LPIPS	TMQI
DSLR	0.500	17.358	0.412	0.761
TBEFN	0.496	13.871	0.379	0.749
Zero-DCE++	0.472	15.761	0.396	0.746
FFDNet + GC	0.526	16.989	0.459	0.680
DFTL (Ours)	<u>0.668</u>	<u>19.905</u>	<u>0.349</u>	<u>0.759</u>
TFDL (Ours)	0.698	20.168	0.293	0.761

Table 2: Quantitative comparison on HDR+ testing set. The best results are printed in boldface. The second best results are underlined.

Methods	SSIM	PSNR	LPIPS	TMQI
DSLR	0.674	14.541	0.388	0.718
TBEFN	-	-	-	-
Zero-DCE++	-	-	-	-
FFDNet + GC	<u>0.542</u>	<u>12.721</u>	0.475	0.686
DFTL (Ours)	0.517	11.727	<u>0.352</u>	<u>0.767</u>
TFDL (Ours)	0.502	11.491	0.329	0.778

Table 3: Quantitative comparison on the LOL dataset. The best results are printed in boldface. The second best results are underlined. Since TBEFN and Zero-DCE++ are trained on this dataset, we remove them from this comparison.

Methods	SSIM	PSNR	LPIPS	TMQI
DSLR	0.529	<u>13.094</u>	<u>0.385</u>	0.706
TBEFN	-	-	-	-
Zero-DCE++	-	-	-	-
FFDNet + GC	<u>0.524</u>	13.656	0.450	0.531
DFTL (Ours)	0.529	12.347	0.383	<u>0.665</u>
TFDL (Ours)	0.517	12.210	<u>0.385</u>	0.646

Table 4: Quantitative comparison on part 1 of the SCIE dataset. The best results are printed in boldface. The second best results are underlined. Since TBEFN and Zero-DCE++ are trained on this dataset, we remove them from this comparison.

For qualitative comparison, examples from the HDR+ testing set with zoomed-in views are shown in Figure 2. We also show examples from the LOL dataset and part 1 of the SCIE dataset in Figure 3 to demonstrate the effectiveness of our methods on an unseen dataset. From these representative examples, we can see that both of our methods outperform the other methods in the comparison in terms of both tone-mapping and denoising. Compared to TBEFN and Zero-DCE++, DSLR does a better job of tone-mapping the image to a natural brightness level, especially in the




Figure 4: Qualitative comparisons between methods on examples from the MEF dataset (top two rows) and LIME dataset (bottom two rows). Four sub-regions selected from each example are zoomed in 5 times to show more details.

extreme low-light scenes. However, all these three methods suffer from the amount of noise in the input. In terms of denoising, FFDNet does a pretty good job at removing heavy noise in the input, but it may overly blur the image and smooth out important details or high-frequency information. Our methods significantly outperform other methods on both the HDR+ dataset and the LOL dataset, in terms

of both the tone-mapping quality and denoising quality. To be specific, our methods can tone-map the low-light inputs to appropriate brightness levels while maintaining natural contrasts. Moreover, our methods effectively remove the noise while keeping most of the important details intact.

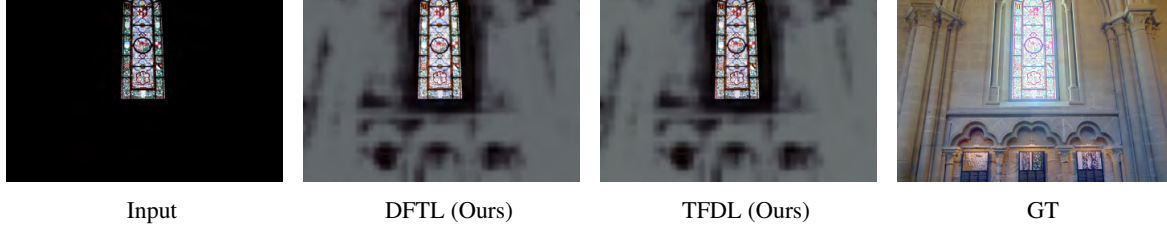


Figure 5: An example of failure cases from the SCIE dataset.

Methods	MACs(G)	Params(M)
DSLR	5.876	14.931
TBEFN	3.682	0.683
Zero-DCE++	0.530	0.011
FFDNet + GC	10.721	0.855
DFTL (Ours)	10.582	39.806
TFDL (Ours)	10.582	39.806

Table 5: Comparison on model sizes and computational complexity. Smallest values are labeled in blue, and greatest values are labeled in red.

4.2.2 No-Reference Comparison

To further showcase the effectiveness of our methods on unseen datasets, we also tested our models on two more popular datasets, namely the MEF dataset [34] and the LIME dataset [18]. Since these datasets do not contain paired ground-truth images, we only show outputs from our models as well as the other methods in Figure 4 for a qualitative comparison. From these no-reference results, we can see that our results have better contrast, look more natural, and are aesthetically more pleasing, compared to the results from other methods.

4.2.3 Computational Complexity Comparison

In order to compare the efficiency of these methods, we also measured the model size and computational complexity of each model and summarized the results in Table 5. To be specific, we measured the number of MAC (multiply-accumulate operations) as well as learnable parameters in these models with a 3-channel input image in size 224×224 . We can see from the table that, although our joint framework performs well in both the quantitative and qualitative comparisons, it does have a trade-off in its model size. However, considering that our method performs both tone-mapping and denoising at the same time and that denoising algorithms are often computationally intensive, this trade-off should be worthy.

4.2.4 Failure Cases

Although our models performed well in most cases, there can be failure cases in some extreme scenarios that are not presented in our HDR+ training dataset. As shown in Figure 5, our models failed in this example from the SCIE dataset since the contrast is too strong and the background is too dark. This is likely due to the lack of similar examples in the HDR+ dataset and the fact that we did not train or fine-tune our models on the samples from the SCIE dataset.

5. Conclusions

In this paper, we have proposed a composite HDR image enhancement model that jointly performs HDR image tone-mapping and image denoising in a multi-scale framework. Based on the testing results on the HDR+ dataset shown in the previous section, our joint multi-scale tone mapping and denoising model for HDR image enhancement outperforms other state-of-the-art image enhancement methods both quantitatively and qualitatively. Our experiments on the ordering of the tone-mapping operator (TMO) and the denoising operator demonstrate that applying tone mapping before denoising yields slightly better results in our joint image enhancement framework. In the future, we plan to exploit the combination of multi-image tone-mapping and multi-image denoising for HDR images.

References

- [1] Mahmoud Afifi, Konstantinos G Derpanis, Björn Ommer, and Michael S Brown. Learning multi-scale photo exposure correction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2021.
- [2] Saeed Anwar, Cong Phuoc Huynh, and Fatih Porikli. Identity enhanced residual image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [3] Long Bao, Zengli Yang, Shuangquan Wang, Dongwon Bai, and Jungwon Lee. Real image denoising based on multi-scale residual dense block and cascaded u-net with block-connection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.

- [4] Wilhelm Burger and Mark J. Burge. *Scale-Invariant Feature Transform (SIFT)*, pages 609–664. Springer London, London, 2016.
- [5] Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018.
- [6] Chen Chen, Qifeng Chen, Minh N. Do, and Vladlen Koltun. Seeing motion in the dark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [7] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [8] Jingwen Chen, Jiawei Chen, Hongyang Chao, and Ming Yang. Image blind denoising with generative adversarial network based noise modeling. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3155–3164, 2018.
- [9] Emily Denton, Soumith Chintala, Arthur Szlam, and Rob Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS’15*, page 1486–1494, Cambridge, MA, USA, 2015. MIT Press.
- [10] Frédo Durand and Julie Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In *Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 257–266, 2002.
- [11] Thibaud Ehret, Axel Davy, Pablo Arias, and Gabriele Facciolo. Joint demosaicing and denoising by overfitting of bursts of raw images. *CoRR*, abs/1905.05092, 2019.
- [12] Minhao Fan, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Integrating semantic segmentation and retinex model for low-light image enhancement. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM ’20, page 2317–2325, New York, NY, USA, 2020. Association for Computing Machinery.
- [13] Michaël Gharbi, Gaurav Chaurasia, Sylvain Paris, and Frédo Durand. Deep joint demosaicking and denoising. *ACM Trans. Graph.*, 35(6), Nov. 2016.
- [14] Clément Godard, Kevin Matzen, and Matt Uyttendaele. Deep burst denoising. *CoRR*, abs/1712.05790, 2017.
- [15] Miguel Granados, Tunç Ozan Aydın, J. Rafael Tena, Jean-François Lalonde, and Christian Theobalt. HDR image noise estimation for denoising tone mapped images. In *Proceedings of the 12th European Conference on Visual Media Production, CVMP ’15*, New York, NY, USA, 2015. Association for Computing Machinery.
- [16] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [17] Shi Guo, Zifei Yan, Kai Zhang, Wangmeng Zuo, and Lei Zhang. Toward convolutional blind denoising of real photographs. *2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [18] Xiaojie Guo, Yu Li, and Haibin Ling. LIME: Low-light image enhancement via illumination map estimation. *IEEE Transactions on Image Processing*, 26(2):982–993, 2017.
- [19] 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 Transactions on Graphics (Proc. SIGGRAPH Asia)*, 35(6), 2016.
- [20] Jingwen He, Yihao Liu, Yu Qiao, and Chao Dong. Conditional sequential modulation for efficient global image retouching. *arXiv preprint arXiv:2009.10390*, 2020.
- [21] Haiyang Jiang and Yinqiang Zheng. Learning to see moving objects in the dark. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [22] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlighten: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 30:2340–2349, 2021.
- [23] Filippos Kokkinos and Stamatis Lefkimmiatis. Deep image demosaicking using a cascade of convolutional residual denoising networks. *CoRR*, abs/1803.05215, 2018.
- [24] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5835–5843, 2017.
- [25] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2Noise: Learning image restoration without clean data. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2965–2974. PMLR, 10–15 Jul 2018.
- [26] Chongyi Li, Chunle Guo, and Chen Change Loy. Learning to enhance low-light image via zero-reference deep curve estimation. In *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [27] Chongyi Li, Jichang Guo, Fatih Porikli, and Yanwei Pang. Lightnet: A convolutional neural network for weakly illuminated image enhancement. *Pattern Recognition Letters*, 104:15–22, 2018.
- [28] Seokjae Lim and Wonjun Kim. DSLR: Deep stacked Laplacian restorer for low-light image enhancement. *IEEE Transactions on Multimedia*, pages 1–1, 2020.
- [29] Lin Liu, Xu Jia, Jianzhuang Liu, and Qi Tian. Joint demosaicing and denoising with self guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [30] Wei Liu, Qiong Yan, and Yuzhi Zhao. Densely self-guided wavelet network for image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.

- [31] Yang Liu, Saeed Anwar, Liang Zheng, and Qi Tian. GradNet image denoising. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [32] Kun Lu and Lihong Zhang. TBEFN: A two-branch exposure-fusion network for low-light image enhancement. *IEEE Transactions on Multimedia*, pages 1–1, 2020.
- [33] Feifan Lv, Bo Liu, and Feng Lu. Fast enhancement for non-uniform illumination images using light-weight CNNs. In *Proceedings of the 28th ACM International Conference on Multimedia*, MM '20, page 1450–1458, New York, NY, USA, 2020. Association for Computing Machinery.
- [34] Kede Ma, Kai Zeng, and Zhou Wang. Perceptual quality assessment for multi-exposure image fusion. *IEEE Transactions on Image Processing*, 24(11):3345–3356, 2015.
- [35] Antoine Monod, Julie Delon, and Thomas Veit. An Analysis and Implementation of the HDR+ Burst Denoising Method. *Image Processing On Line*, 11:142–169, 2021. <https://doi.org/10.5201/ipol.2021.336>.
- [36] Nicola Pierazzo, Jean-Michel Morel, and Gabriele Facciolo. Multi-Scale DCT Denoising. *Image Processing On Line*, 7:288–308, 2017. <https://doi.org/10.5201/ipol.2017.201>.
- [37] Wenqi Ren, Sifei Liu, Lin Ma, Qianqian Xu, Xiangyu Xu, Xiaochun Cao, Junping Du, and Ming-Hsuan Yang. Low-light image enhancement via a deep hybrid network. *IEEE Transactions on Image Processing*, 28(9):4364–4375, 2019.
- [38] Chunwei Tian, Yong Xu, Lunke Fei, and Ke Yan. Deep learning for image denoising: A survey. In Jeng-Shyang Pan, Jerry Chun-Wei Lin, Bixia Sui, and Shih-Pang Tseng, editors, *Genetic and Evolutionary Computing*, pages 563–572, Singapore, 2019. Springer Singapore.
- [39] Danai Triantafyllidou, Sean Moran, Steven McDonagh, Sarah Parisot, and Gregory Slabaugh. Low light video enhancement using synthetic data produced with an intermediate domain mapping. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 103–119, Cham, 2020. Springer International Publishing.
- [40] Gregory Vaksman, Michael Elad, and Peyman Milanfar. Lidia: Lightweight learned image denoising with instance adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2020.
- [41] Li-Wen Wang, Zhi-Song Liu, Wan-Chi Siu, and Daniel P. K. Lun. Lightening network for low-light image enhancement. *IEEE Transactions on Image Processing*, 29:7984–7996, 2020.
- [42] Ruixing Wang, Qing Zhang, Chi-Wing Fu, Xiaoyong Shen, Wei-Shi Zheng, and Jiaya Jia. Underexposed photo enhancement using deep illumination estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [43] Yang Wang, Yang Cao, Zheng-Jun Zha, Jing Zhang, Zhiwei Xiong, Wei Zhang, and Feng Wu. Progressive retinex: Mutually reinforced illumination-noise perception network for low-light image enhancement. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, page 2015–2023, New York, NY, USA, 2019. Association for Computing Machinery.
- [44] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference*, 2018.
- [45] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *BMVC*, 2018.
- [46] Ke Xu, Xin Yang, Baocai Yin, and Rynson W.H. Lau. Learning to restore low-light images via decomposition-and-enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [47] Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, and Jiaying Liu. From fidelity to perceptual quality: A semi-supervised approach for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [48] Hojatollah Yeganeh and Zhou Wang. Objective quality assessment of tone-mapped images. *IEEE Transactions on Image Processing*, 22(2):657–667, 2013.
- [49] Guoshen Yu and Guillermo Sapiro. DCT image denoising: a simple and effective image denoising algorithm. *Image Processing On Line*, 1, 2011.
- [50] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing*, 26(7):3142–3155, 2017.
- [51] Kai Zhang, Wangmeng Zuo, and Lei Zhang. FFDNet: Toward a fast and flexible solution for CNN-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018.
- [52] Lin Zhang, Lijun Zhang, Xiao Liu, Ying Shen, Shaoming Zhang, and Shengjie Zhao. Zero-shot restoration of backlit images using deep internal learning. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, page 1623–1631, New York, NY, USA, 2019. Association for Computing Machinery.
- [53] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.
- [54] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *Proceedings of the 27th ACM International Conference on Multimedia*, MM '19, page 1632–1640, New York, NY, USA, 2019. Association for Computing Machinery.
- [55] Anqi Zhu, Lin Zhang, Ying Shen, Yong Ma, Shengjie Zhao, and Yicong Zhou. Zero-shot restoration of underexposed images via robust retinex decomposition. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6, 2020.
- [56] Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. EEMEFN: Low-light image enhancement via edge-enhanced multi-exposure fusion network. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(07):13106–13113, Apr. 2020.