

# Towards Fast and Light-Weight Restoration of Dark Images

Mohit Lamba\*

ee18d009@smail.iitm.ac.in

Atul Balaji

ee16b002@smail.iitm.ac.in

Kaushik Mitra

kmitra@ee.iitm.ac.in

Computational Imaging Lab

Dept. of Electrical Engineering

Indian Institute of Technology Madras

Chennai, India

---

## Abstract

The ability to capture good quality images in the dark and *near-zero lux* conditions has been a long-standing pursuit of the computer vision community. The seminal work by Chen *et al.* [5] has especially caused renewed interest in this area, resulting in methods that build on top of their work in a bid to improve the reconstruction. However, for practical utility and deployment of low-light enhancement algorithms on edge devices such as embedded systems, surveillance cameras, autonomous robots and smartphones, the solution must respect additional constraints such as limited GPU memory and processing power. With this in mind, we propose a deep neural network architecture that aims to strike a balance between the network latency, memory utilization, model parameters, and reconstruction quality. The key idea is to forbid computations in the High-Resolution (HR) space and limit them to a Low-Resolution (LR) space. However, doing the bulk of computations in the LR space causes artifacts in the restored image. We thus propose *Pack* and *UnPack* operations, which allow us to effectively transit between the HR and LR spaces without incurring much artifacts in the restored image. State-of-the-art algorithms on dark image enhancement need to pre-amplify the image before processing it. However, they generally use ground truth information to find the amplification factor even during inference, restricting their applicability for unknown scenes. In contrast, we propose a simple yet effective light-weight mechanism for automatically determining the amplification factor from the input image. We show that we can enhance a full resolution,  $2848 \times 4256$ , extremely dark single-image in the ballpark of 3 seconds even on a CPU. We achieve this with  $2 - 7 \times$  fewer model parameters,  $2 - 3 \times$  lower memory utilization,  $5 - 20 \times$  speed up and yet maintain a competitive image reconstruction quality compared to the state-of-the-art algorithms.

## 1 Introduction

The ability to swiftly capture high quality images with modest computations has led to the widespread proliferation of digital images. These advantages are, however, limited to good lighting conditions. Achieving similar results under low light is still a significant challenge. While much of the work in this direction has focused on enhancing weakly illuminated

---

\* Corresponding author

© 2020. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.




Figure 1: Performance comparison of the proposed method with state-of-the-art methods Chen *et al.* [6], Gu *et al.* [10], Maharjan *et al.* [23] and traditional methods LIME [11], and Li *et al.* [20] for extreme low-light single-image enhancement. Refer to Table 1 for more details.

images [16, 18, 20, 27, 28, 40], enhancement of extremely dark images has received comparatively lesser attention.

Recently, however, a landmark paper by Chen *et al.* [6] has shown that it is possible to restore extremely dark images captured under near-zero lux conditions. Following this work, several modifications have been proposed in a bid to improve the reconstruction quality. This includes the incorporation of attention units [10], recurrent units [9], the adoption of a multi-scale approach [10, 24] and the usage of deeper networks [23]. With these added complexities, these methods are constrained to run on desktop GPUs such as NVIDIA RTX 2080Ti with 12GB storage. But, real-world applications require image enhancement algorithms to run on embedded systems and edge devices with limited CPU RAM or minimal GPU capacity. One possible solution is to process the images in VGA resolution [10, 20, 22, 68, 40, 48]. But, this is in contrast to the current trend of capturing and processing high-definition images. Consequently, we aim to design a deep network that can restore an extreme low-light high-definition single-image with minimal CPU latency and low memory footprint, but at the same time has a competitive image restoration quality.

We propose a deep neural network, called Low-Light Packing Network (LLPackNet), which is faster and computationally cheaper than the existing solutions. Recognizing the fact that a neural network’s complexity increases quadratically with spatial dimensions [56], we perform the bulk of computations in a much lower resolution by performing aggressive down/up sampling operation. This is in contrast with much of the existing literature that down/up sample the feature maps in gradations [6, 20, 68, 47, 49], which increases network latency and memory utilization. For performing large downsampling operations, popular choices such as max-pooling and strided convolution [20] cannot be used as they would cause much loss in information. We therefore, propose *Pack*  $\alpha \times$  downsampling operation, which rearranges the pixels in such a manner that it reduces the spatial dimension by a factor of  $\alpha$ , while increasing the number of channels by a factor of  $\alpha^2$ , see Fig. 2. We show that the *Pack* operation bestows LLPackNet with an enormous receptive field which is not trivially possible by directly operating in the HR space. We also propose *UnPack*  $\alpha \times$  operation, which complements the *Pack*  $\alpha \times$  operation to do large upsampling. This operation is much faster than the usual transposed convolution layer [20] and has no learnable parameters. For upsampling, PixelShuffle [63] is another viable option but it lacks proper correlation between the color channels and hence results in heavy color cast in the restored image as shown in Fig. 5. Altogether, the proposed *Pack* and *UnPack* operations allow us to operate in a much lower resolution space for computational advantages, without significantly affecting the restoration

quality. See Fig. 1 for a qualitative comparison with state-of-the-art algorithms.

State-of-the-art deep learning solutions on extreme low-light image enhancement need to pre-amplify dark images before processing them [10, 9, 11, 23]. However, these methods use ground-truth knowledge for predicting the amplification factor. In a real-world setting, because of lack of ground-truth (GT) knowledge, the amplification factor cannot be estimated properly and hence this would lead to degradation in performance. We therefore, equip the proposed LLPackNet with an amplifier module, which will estimate the amplification factor directly from the input image histogram.

To summarize, the main contributions of this paper are as follows — 1) We propose a deep neural network architecture, called *LLPackNet*, that enhances an extremely dark single-image at high resolution even on a CPU with very low latency and computational resources. 2) We propose *Pack* and *UnPack* operations for better color restoration. 3) LLPackNet can estimate the amplification factor directly from the input image, without relying on ground-truth information, making it practical for real world applications. 4) Our experiments show that compared to existing solutions, we are able to restore high definition, extreme low-light RAW images with  $2\text{--}7\times$  fewer model parameters,  $2\text{--}3\times$  lower memory and  $5\text{--}20\times$  speed up, with a competitive restoration quality. Our code is available at <https://github.com/MohitLamba94/LLPackNet>.

## 2 Related Work

Low-light enhancement methods are chiefly comprised of histogram equalization [16, 21, 28], Retinex based decomposition [9, 11, 18, 20, 25, 45] and Deep learning based methods [6, 7, 19, 22, 30, 32, 37, 39, 40, 46]. Most of them however, do not target extreme low-light conditions or high resolution images. More recently, Chen *et al.* [5] proposed an end-to-end pipeline to restore extreme low-light high-definition RAW images, which has spurred several other works in this direction [1, 3, 11, 15, 23, 24]. Most of these methods, however, have significantly large processing time and memory utilization. As noted in Sec. 1 many of them also require GT information for image pre-amplification. However, other image amplification techniques that involve the use of CRF [31, 44], image histogram [16, 28] or other assumptions [11, 43] have been used in traditional image enhancement methods to estimate amplification, using only the input image. Borrowing from these ideas, we develop an amplifier module that uses the histogram of the input dark image to predict the amplification factor automatically, without relying on GT information. To the best of our knowledge, this has not been attempted before for deep learning based dark image enhancement.

Fast and efficient CNN models have been explored in other areas, especially image classification, but is mostly achieved by either approximating [29, 42] or pruning the learned weights [8]. In contrast, we propose a network that is inherently fast and efficient without using such weight-approximation or pruning approaches.

## 3 Low-Light Packing Network (LLPackNet)

We propose Low-Light Packing Network (LLPackNet) for enhancing extremely dark high resolution single-images with low time-memory complexity. We first describe the network architecture, shown Fig. 2, in Sec. 3.1 and then analyze the important components of our network, the Pack and UnPack operations, in Sec. 3.2.




Figure 2: Our proposed network, LLPackNet, learns the required amplification factor directly from the dark input image and uses the novel *Pack* and *UnPack* operations to perform aggressive down/up sampling with minimal color distortions. The network has low latency and low memory footprint.

### 3.1 Network architecture

**Image amplification:** In general, dark images need to be pre-amplified before enhancing them. We estimate the amplification factor using the incoming RAW image  $I_{i/p}^{HR}$  by constructing a 64 bin histogram, with the histogram bins being equidistant in the log domain. This provides a finer binning resolution for lower intensities and a coarser binning resolution for higher intensities. The histogram is used by a multilayer perceptron, having just one hidden layer, to estimate the amplification factor.

**Fast and light-weight enhancement:** As discussed in Sec 1, we want to perform most of the processing in LR space. Hence, our first step is to downsample the input image, without losing any information. For this purpose, we propose *Pack*  $\alpha \times$  operation, that downsamples the image by a factor of  $\alpha$  along the spatial dimensions while increasing the number of channels by a factor of  $\alpha^2$ . This is shown in Fig. 2 (b) for  $\alpha = 2$ . A pseudo code is also provided in Algorithm 1. Our goal is to perform  $16 \times$  downsampling, which we do in two stages. In the first stage, the *Pack*  $2 \times$  operation separates out the red, green and blue color components lying in the  $2 \times 2$  Bayer pattern [13] of the amplified image  $I_{i/p}^{HR}$ . This reduces the spatial dimension by half and increases the channels from 1 to 4 ( $2^2$ ). Once the colors are separated into these channels, a subsequent *Pack*  $8 \times$  operation is applied individually on each color channel, further reducing the spatial dimension from  $2 \times$  to  $16 \times$  lower resolution but increasing the number of channels from 1 to 64 ( $8^2$ ). Now, using a  $3 \times 3$  convolution kernel, the channel dimension of each color component is reduced such that on concatenation, the resulting feature map has only 60 channels. The channel reduction at this stage is essential to prevent parameter and memory explosion in the downstream operations. This downsampled representation is then processed by a series of convolution operations. For this purpose, we use the Residual Dense Network [49] (RDN) — which consists of 3 residual dense blocks each with 6 convolutional layers and a growth rate of 32. RDN does not perform any down/up sampling operation or cause any change in channel dimension in its output. The output of the RDN now needs to be upsampled and for this we use the proposed *UnPack*  $2 \times$  operation, which is the inverse of *Pack*  $2 \times$ . *UnPack*  $2 \times$ , however, reduces the number of channels from 60 to 15 ( $60/2^2$ ) and this needs to be increased to 192 ( $8^2 \times 3$ ) to allow the final  $8 \times$  upsampling using *UnPack*  $8 \times$ . For this we use another set of  $3 \times 3$  convolutions. Except for this operation, all the computations are done in the  $16 \times$  lower resolution. We finally perform *UnPack*  $8 \times$  operation to get the restored image,  $I_{o/p}^{HR}$ .

**Loss function:** Similar to Ignatov *et al.* [14], we compute the color loss, content loss and total variation (TV) loss on the restored image to train the network. Specifically, we use,

$$\begin{aligned}
 Loss = & \lambda_1 \cdot \|GT - I_{o/p}^{HR}\|_1 + \lambda_2 \cdot \|\Psi(GT) - \Psi(I_{o/p}^{HR})\|_1 + \lambda_3 \cdot \|\delta(GT) - \delta(I_{o/p}^{HR})\|_1 \\
 & + \lambda_4 \cdot TV(I_{o/p}^{HR}) + \lambda_5 \cdot \|w\|_1
 \end{aligned} \tag{1}$$

where  $\Psi$  is a feature map of VGG-19,  $\delta$  performs Gaussian smoothing and  $w$  denotes the network weights. VGG-19 features are obtained right after the final 3 max-pool layers.

## 3.2 Pack/UnPack operation for better color restoration

The last section discussed LLPackNet from the vantage point of network complexity. In this section we analyze the network from the standpoint of reconstruction quality.

---

**Algorithm 1:** Python code for performing Pack  $\alpha \times$  and UnPack  $\alpha \times$  operation.

---

### *Pack $\alpha \times$ operation*

**Input:** An RGB image  $I_{HR}$  of dimensions  $H \times W \times 3$ .

**Output:**  $I_{LR}$  of dimension  $\frac{H}{\alpha} \times \frac{W}{\alpha} \times 3\alpha^2$ .

*count* = 0

**for** *row* **in** range( $\alpha$ ) :

**for** *col* **in** range( $\alpha$ ) :

$I_{LR}[:, :, \textit{count} : \textit{count} + 3] = I_{HR}[\textit{row} : H : \alpha, \textit{col} : W : \alpha, :]$

*count* = *count* + 3

### *UnPack $\alpha \times$ operation*

**Input:**  $I_{LR}$  of dimension  $\frac{H}{\alpha} \times \frac{W}{\alpha} \times 3\alpha^2$ .

**Output:** An RGB image  $I_{HR}$  of dimensions  $H \times W \times 3$ .

*count* = 0

**for** *row* **in** range( $\alpha$ ) :

**for** *col* **in** range( $\alpha$ ) :

$I_{HR}[\textit{row} : H : \alpha, \textit{col} : W : \alpha, :] = I_{LR}[:, :, \textit{count} : \textit{count} + 3]$

*count* = *count* + 3

---

**Improving color correlation with UnPack  $\alpha \times$ :** Making abrupt transitions between LR and HR spaces introduces several distortions in the restored image. To minimize these, we propose the novel Pack  $\alpha \times$  and UnPack  $\alpha \times$  operations. To understand these operations, it is crucial to analyze PixelShuffle [33] - a fast and effective upsampling method, based on which they are formulated. Using an analysis similar to [1, 33, 34], we will show that Pack/UnPack operations lead to better color correlation than PixelShuffle.

First we analyze the PixelShuffle operation. In Fig. 3 a),  $T^{LR}$  refers to the penultimate feature map, which is upsampled with zero padding and then convolved with  $w^{HR}$  to obtain the restored image  $O^{HR}$ . We now explain the color coding used in the figure. When  $w^{HR}$  convolves with  $T^{HR}$ , for each shifted position of  $w^{HR}$ , only the weights in one set of colors in  $w^{HR}$  contribute to an output pixel in  $T^{HR}$ . We label the output pixel with the same color. Doing convolution in HR is computationally expensive. However, an equivalent operation in the LR space can be performed as shown in Fig. 3 b). This involves decomposing  $w^{HR}$  into smaller kernels of  $w^{LR}$  which are then convolved with  $T^{LR}$  to produce  $O^{LR}$ . Using PixelShuffle,  $O^{HR}$  can then be obtained from  $O^{LR}$ . However, in this scheme, each kernel in  $w^{HR}$  maintains a monopoly on one of the red, green or blue color channels in the restored image  $O^{HR}$ , see Fig. 3 c). Thus, restoring images using PixelShuffle causes weak correlation among the color channels of  $O^{HR}$ , leading to color artifacts as shown in Fig. 5.

The goal of UnPack operation is to enhance the correlation among the color channels of  $O^{HR}$ . For this purpose, along with the upsampling, zero-padding and convolution operations, we introduce a re-grouping step as shown in Fig. 3 d). This may appear to be a complicated two-stage operation, but using our UnPack operation we can easily perform an equivalent

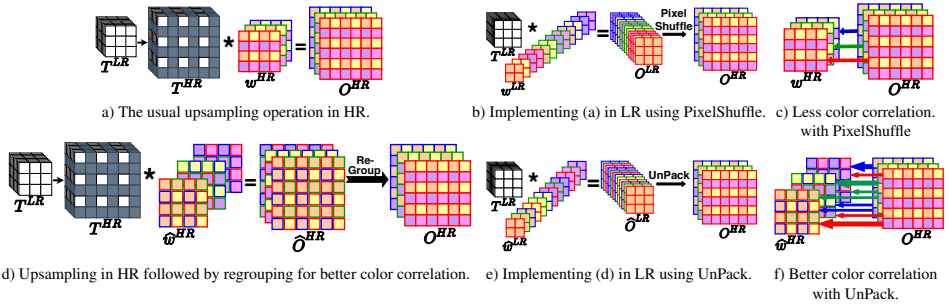


Figure 3: Performing aggressive down/up-sampling causes several color distortions in the restored image. The proposed UnPack operation limits this by improving the color correlation in the restored image and simultaneously performing quick upsampling from LR to HR space. The effectiveness of the proposed solution is demonstrated in Fig. 5.

operation in the LR space, as shown in Fig. 3 e). Note that this operation has the same time complexity as the operation shown in Fig. 3 b). For this operation, we decompose  $\hat{w}_{HR}$  into  $\hat{w}_{LR}$  and then apply UnPack  $\alpha \times$ . From Fig. 3 f), we see that all the kernels of  $\hat{w}_{HR}$  are collectively responsible for all the colors in  $O_{HR}$ . Thus, UnPack operation leads to better color correlation than PixelShuffle.

The effectiveness of the proposed UnPack operation can also be intuitively understood in the LR space by comparing Fig. 3 b) and Fig. 3 e). UnPack preserves the RGB ordering in the LR, whereas, PixelShuffle breaks this ordering, especially for large upsampling factors. For example, for a given spatial location in HR, PixelShuffle separates the Red and Blue pixels by  $2 \times 8^2 = 128$  channels in LR for  $8 \times$  upsampling. UnPack, however, always separates them by only 1 Green pixel for any upsampling factor. This is crucial because for CNNs it is well known that nearby features correlate more than spaced out ones [65]. Thus even though, UnPack does not introduce any new parametrization, its arrangement favors better color restoration. Therefore, in Fig. 5, PixelShuffle’s restored image is heavily affected by color cast, but no such distortion is observed in the case of UnPack.

**Increasing receptive field with Pack  $\alpha \times$ :** Having a large receptive field is essential for capturing the contextual information in an image. Downsampling the incoming feature map using the novel Pack operation equips LLPackNet with a large receptive field. To illustrate this fact, let us consider a large feature map  $I^{HR}$  which is downsampled to  $I^{LR}$  using Pack  $10 \times$  operation. Note that the neighboring pixels in  $I^{LR}$  are actually 10 pixels apart in  $I^{HR}$ . Also, the pixels along the channel dimension of  $I^{LR}$  are in a  $10 \times 10$  neighborhood in  $I^{HR}$ . Thus, even using a  $3 \times 3$  convolution kernel on  $I^{LR}$  with a stride of 1 leads to a receptive field of 900 pixels in  $I^{HR}$ . In contrast, to do a similar operation directly on  $I^{HR}$ , requires a  $30 \times 30$  kernel with a stride of 10, which is impractical.

## 4 Experiments

### 4.1 Experimental settings

For extreme low-light single-image enhancement, we compare with Chen *et al.* [9], Gu *et al.* [10] and Maharjan *et al.* [23]. In addition, we also tried conventional techniques such as LIME [11] and Li *et al.* [20] but they did not work well for dark images. The publicly available training and test codes of these methods have been used for the comparisons.

Model	Processing Time (in seconds)	Memory (in GB)	Parameters (in million)	PSNR(dB) / SSIM	
				w/o GT exposure	using GT exposure
Maharjan <i>et al.</i> [23]	120	10	2.5	20.98 / 0.49	28.41 / <b>0.81</b>
Gu <i>et al.</i> [10]	77	8	3.5	21.90 / 0.59	<b>28.53 / 0.81</b>
Chen <i>et al.</i> [5]	17	5	7.75	22.93 / 0.70	28.30 / 0.79
Chen <i>et al.</i> [5] + Our Amplifier	17	5	7.76	22.98 / <b>0.71</b>	28.30 / 0.79
LLPackNet (Ours)	3	3	<b>1.1</b>	<b>23.27</b> / 0.69	27.83 / 0.75

Table 1: Results on the SID dataset [5] for extremely dark  $2848 \times 4256$  RAW images. Compared to existing approaches, we have 2–7× fewer model parameters, 2–3× lower memory, 5–20× speed up with competitive restoration quality.

For experiments on dark images, we use See-in-the-Dark (SID) dataset [5] captured with high definition full-frame Sony  $\alpha 7S$  II Bayer sensor. Unlike some methods that collect their dataset by simulating pairs of low-light and GT images [19, 22, 25, 30, 32, 40], SID provides physically captured extreme low-light RAW images of resolution  $2848 \times 4256$ . We additionally show comparisons on the LOL dataset [40] to evaluate the performance of LLPackNet on a notably distinct test set-up. In contrast to SID, LOL has weakly illuminated VGA resolution PNG compressed images. Additionally, SID comes with GT and low-light exposure information, which can be used for estimating the pre-amplification factor, but LOL has no such information.

We use the train/test split as given in the datasets. For LLPackNet, patches of size  $512 \times 512$  are used for training and full resolution for testing. For benchmarking, we use the PyTorch [26] framework on Intel Xeon E5-1620V4 @ 3.50 GHz CPU with 64 GB RAM. We use the default Adam optimizer of PyTorch with fixed learning rate of  $10^{-4}$ . All convolutions use kernels of size  $3 \times 3$  with *He initialization* [12]. Our network was allowed to train for 400,000 iterations. We use,  $\lambda_1 = 1$ ,  $\lambda_2 = 3$ ,  $\lambda_3 = 1$ ,  $\lambda_4 = 400$  and  $\lambda_5 = 10^{-6}$ .

## 4.2 Restoration results for extreme low-light images

We compare our network with Chen *et al.* [5], Gu *et al.* [10] and Maharjan *et al.* [23] on the SID dataset, see Table 1 and Fig. 4. These methods use the ratio of GT exposure to that of the input dark image, available in the SID dataset, to pre-amplify the images. The corresponding results are shown under the label ‘using GT exposure’ in Table 1 and Fig. 4. But, since the GT information will not be readily available in a real-world setting, we additionally show results in the absence of GT information. This is shown under the heading ‘w/o GT exposure’. We also show results for ‘Chen *et al.* + Our Amplifier’ in which our proposed amplifier is added to their algorithm. We have chosen Chen *et al.* because they have the least time and memory complexity, compared to the other existing methods. All the methods are appropriately retrained before evaluation.

**Network speed and memory utilization:** As shown in Table 1, LLPackNet is 5–20× faster with 2–3× lower memory and 2–7× lesser model parameters. We achieve this because we do the bulk of operations in  $16 \times$  lower resolution. In contrast, Maharjan *et al.* [23] do not perform any downsampling operation and therefore, the feature maps propagating through their network are huge. This results in very high network latency and memory consumption. Gu *et al.* [10] adopt a multi-scale approach that requires feature map propagation at  $2 \times$  and  $4 \times$  lower resolution. But this marginal downsampling is not sufficient to contain the network latency and memory consumption. Chen *et al.* [5] have relatively better metrics by performing up to  $32 \times$  downsampling. But this is done only in steps of 2, requiring five downsampling and five upsampling operations. Further, four out of five upsampling oper-



Figure 4: [View with good screen brightness] Performance comparison of the proposed LLPackNet with state-of-the-art algorithms corresponding to Table 1. (A): All the methods work well with GT exposure. (B): But in a realistic scenario, where exact GT exposure may not be available during inference, only the proposed LLPackNet gives proper restoration.

ations are done using transposed convolution [17], which is much slower than the proposed UnPack operation. Thus, Chen *et al.* have a moderately high processing time and memory utilization. Check the supplementary for more details.

**Restoration quality:** All methods perform notably well when the GT exposure is available. But in a practical setting when GT exposure is not readily available, except for our LLPackNet, the other methods struggle to restore proper colors. The results for this practical setting are also shown in Fig 1. Adding our amplifier module to Chen *et al.* improves their performance to some extent, but the restored images still exhibit noisy patches and color cast. This is because amplification is not the only deciding factor in improving the performance of a network. Rather, having a large receptive field, which provides more contextual information, and better correlation among the color channels is more important than the correct



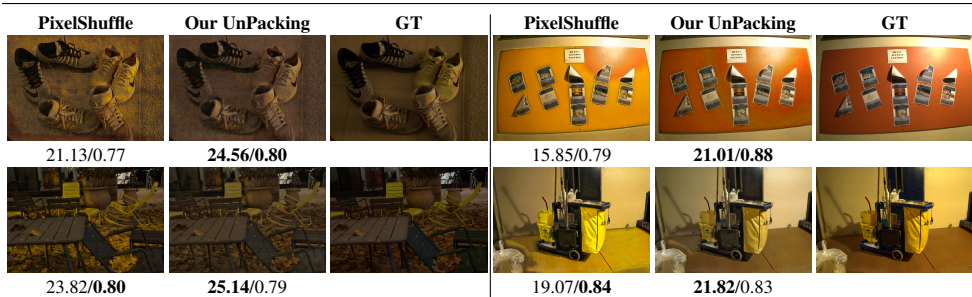


Figure 5: [View with good screen brightness] LLPackNet restoration using PixelShuffle and our proposed UnPack operation. The UnPack operation helps achieve better color restoration by containing color distortions such as color cast.

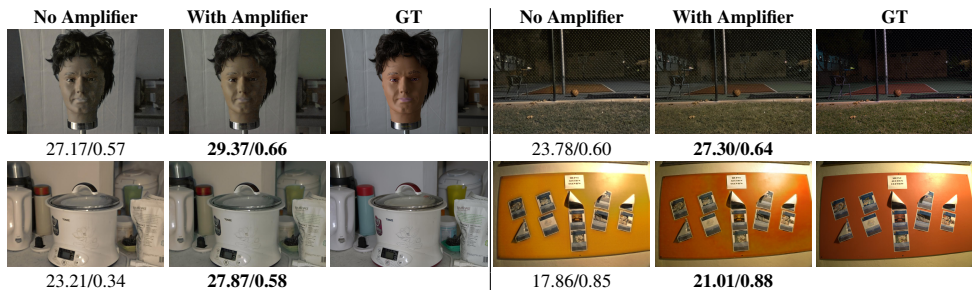


Figure 6: [View with good screen brightness] Results using the proposed LLPackNet with and without amplification estimation. Without amplification, the colors are pale and tend to be monochromatic.

amplification factor. To further assess these claims, refer to the ablation studies in section 4.3, which show that LLPackNet continues to give structurally consistent results even when the amplifier is removed.

### 4.3 Ablation studies on LLPackNet

We now show ablation studies on LLPackNet to better understand the contribution of individual components. For each ablation study the network is appropriately retrained.

**UnPack vs. PixelShuffle:** As a first ablation study, we replace the UnPack operation in the proposed LLPackNet with the PixelShuffle operation [13] and the results are shown in Fig. 5. We notice that the images restored using PixelShuffle are affected by heavy color cast. Using the UnPack operation in place of PixelShuffle improves the PSNR/SSIM from 22.72 dB/0.68 to 23.27 dB/0.69. Thus, the UnPack operation favors better color restoration.

In Section 3.2 we ascribed the better color restoration performance of Pack/UnPack over PixelShuffle to the fact that PixelShuffle breaks the RGB ordering in LR space, especially for large upsampling factors, whereas UnPack preserves the RGB ordering for any factor. To further test this hypothesis, we conducted an ablation study where we changed the image downsampling factor from  $16\times$  to  $8\times$ , so that the final UnPack/PixelShuffle operation performs  $4\times$  upsampling instead of  $8\times$ . This results in increased time and computational complexity for LLPackNet, but it reduces the separation between the Red and Blue channels in PixelShuffle. The performance of UnPack (23.29 dB) is almost the same as the  $16\times$  upsampling case, however the performance of PixelShuffle (23.28 dB) improves. In

Model	Processing Time	PSNR (dB)	SSIM
Chen <i>et al.</i> [8]	0.21 sec.	18.82	0.73
LIME [14]	0.19 sec.	16.94	0.60
Li <i>et al.</i> [14]	17.89 sec.	13.85	0.65
Gu <i>et al.</i> [14]	0.41 sec.	19.46	<b>0.75</b>
LLPackNet-8 $\times$ (Proposed)	<b>0.06</b> sec.	<b>19.61</b>	0.69
LLPackNet-4 $\times$ (Proposed)	0.24 sec.	19.60	0.74

Table 2: Results on the LOL dataset [14] for weakly illuminated compressed images having  $400 \times 600$  VGA resolution. LLPackNet with  $8\times$  downsampling (LLPackNet- $8\times$ ) is very fast. But, since the image resolution is quite low, LLPackNet- $4\times$  opts for smaller downsampling to achieve better reconstruction as reflected in the SSIM value.

the case of  $8\times$  upsampling UnPack shows a gain of about 0.6 dB over PixelShuffle but for the case of  $4\times$  upsampling the difference reduces to 0.01dB. This confirms our hypothesis that Pack/UnPack has better performance because it preserves the RGB ordering for any upsampling factor.

**Estimating proper amplification:** Fig. 6 shows the restoration results using LLPackNet with and without the amplifier. Similar to the scotopic vision [9, 14], without the amplifier, the restoration has faded colors. But the removal of the amplifier does not induce the annoying artifacts seen in restoration done using Chen *et al.* (see Fig. 4 (B)). This can be attributed to the large receptive field due to the Pack operation. With the amplifier the performance of the network improves from 22.53dB/0.66 to 23.27dB/0.69.

Overall the combined effect of using the proposed Pack/UnPack operation over PixelShuffle, and estimating proper amplification, increases the average PSNR/SSIM from 21.35 dB / 0.60 to 23.27 dB / 0.69.

#### 4.4 LLPackNet for low-resolution images

The SID dataset contains high definition images, thereby, allowing us to chose a large downsampling factor of 16. This leads us to the question: Can LLPackNet also work for LR images? When LR images are downsampled using a large factor, the intra-channel correlation in the downsampled image is reduced, which negatively impacts the restoration. To investigate this, we conducted experiments on the LOL dataset [14] containing weakly illuminated images at VGA resolution of  $400 \times 600$ . As the images in the LOL dataset are already in the compressed PNG format, the  $2\times$  downsampling at the beginning of LLPackNet to separate out the Bayer pattern is not required. Thus, the effective downsampling is only  $8\times$  and we denote this network by LLPackNet- $8\times$ . The results are shown in Table 2. Once again, LLPackNet has the lowest processing time. We further observe that the large receptive field of LLPackNet enhances the denoising and color restoration capabilities. But, a slight blur is also introduced. To verify that the blur is because of large downsampling, we retrain LLPackNet on the LOL dataset with  $4\times$  downsampling, which we denote as LLPackNet- $4\times$ . With LLPackNet- $4\times$ , we obtain sharper results having higher SSIM values.

## 5 Conclusion

In this paper, a fast and light-weight extreme low-light image enhancement network (LLPackNet) has been presented. LLPackNet performs aggressive down/up-sampling using the proposed Pack/UnPack operations to obtain a large receptive field and better color restoration. The network also uses a novel amplifier module that amplifies the input image without relying on ground-truth information. Overall, LLPackNet is  $5\text{--}20\times$  faster and  $2\text{--}3\times$  lighter, and yet maintains a competitive restoration quality compared to state-of-the-art algorithms.

## References

- [1] Sophy Ai and Jangwoo Kwon. Extreme low-light image enhancement for surveillance cameras using attention u-net. *Sensors*, 20(2):495, 2020.
- [2] Andrew Aitken, Christian Ledig, Lucas Theis, Jose Caballero, Zehan Wang, and Wenzhe Shi. Checkerboard artifact free sub-pixel convolution: A note on sub-pixel convolution, resize convolution and convolution resize. *arXiv preprint arXiv:1707.02937*, 2017.
- [3] Yuantian Cai and U Kintak. Low-light image enhancement based on modified u-net. In *2019 International Conference on Wavelet Analysis and Pattern Recognition (ICWAPR)*, pages 1–7. IEEE, 2019.
- [4] Bo Chen and Pietro Perona. Seeing into darkness: Scotopic visual recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3826–3835, 2017.
- [5] 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*, pages 3291–3300, 2018.
- [6] Yu Cheng, Jia Yan, and Zhou Wang. Enhancement of weakly illuminated images by deep fusion networks. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 924–928. IEEE, 2019.
- [7] Vincent Dumoulin and Francesco Visin. A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*, 2016.
- [8] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019.
- [9] Sanjay Ghosh and Kunal N Chaudhury. Fast bright-pass bilateral filtering for low-light enhancement. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 205–209. IEEE, 2019.
- [10] Shuhang Gu, Yawei Li, Luc Van Gool, and Radu Timofte. Self-guided network for fast image denoising. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2511–2520, 2019.
- [11] 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, 2016.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [13] Keigo Hirakawa and Thomas W Parks. Adaptive homogeneity-directed demosaicing algorithm. *IEEE Transactions on Image Processing*, 14(3):360–369, 2005.

- [14] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3277–3285, 2017.
- [15] Tomas Jenicek and Ondrej Chum. No fear of the dark: Image retrieval under varying illumination conditions. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9696–9704, 2019.
- [16] Yeong-Taeg Kim. Contrast enhancement using brightness preserving bi-histogram equalization. *IEEE transactions on Consumer Electronics*, 43(1):1–8, 1997.
- [17] Hunsang Lee, Kwanghoon Sohn, and Dongbo Min. Unsupervised low-light image enhancement using bright channel prior. *IEEE Signal Processing Letters*, 27:251–255, 2020.
- [18] Bo Li, Shuhang Wang, and Yanbing Geng. Image enhancement based on retinex and lightness decomposition. In *2011 18th IEEE International Conference on Image Processing*, pages 3417–3420. IEEE, 2011.
- [19] 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.
- [20] Mading Li, Jiaying Liu, Wenhan Yang, Xiaoyan Sun, and Zongming Guo. Structure-revealing low-light image enhancement via robust retinex model. *IEEE Transactions on Image Processing*, 27(6):2828–2841, 2018.
- [21] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017.
- [22] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognition*, 61:650–662, 2017.
- [23] Paras Maharjan, Li Li, Zhu Li, Ning Xu, Chongyang Ma, and Yue Li. Improving extreme low-light image denoising via residual learning. In *2019 IEEE International Conference on Multimedia and Expo (ICME)*, pages 916–921. IEEE, 2019.
- [24] Sameer Malik and Rajiv Soundararajan. Llrnet: A multiscale subband learning approach for low light image restoration. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 779–783. IEEE, 2019.
- [25] Seonhee Park, Soohwan Yu, Byeongho Moon, Seungyong Ko, and Joonki Paik. Low-light image enhancement using variational optimization-based retinex model. *IEEE Transactions on Consumer Electronics*, 63(2):178–184, 2017.
- [26] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019.

- [27] Etta D Pisano, Shuquan Zong, Bradley M Hemminger, Marla DeLuca, R Eugene Johnston, Keith Muller, M Patricia Braeuning, and Stephen M Pizer. Contrast limited adaptive histogram equalization image processing to improve the detection of simulated spiculations in dense mammograms. *Journal of Digital imaging*, 11(4):193, 1998.
- [28] Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987.
- [29] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016.
- [30] 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.
- [31] Yurui Ren, Zhenqiang Ying, Thomas H Li, and Ge Li. Lecarm: low-light image enhancement using the camera response model. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(4):968–981, 2018.
- [32] Liang Shen, Zihan Yue, Fan Feng, Quan Chen, Shihao Liu, and Jie Ma. Msr-net: Low-light image enhancement using deep convolutional network. *arXiv preprint arXiv:1711.02488*, 2017.
- [33] Wenzhe Shi, Jose Caballero, Ferenc Huszár, Johannes Totz, Andrew P Aitken, Rob Bishop, Daniel Rueckert, and Zehan Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1874–1883, 2016.
- [34] Wenzhe Shi, Jose Caballero, Lucas Theis, Ferenc Huszar, Andrew Aitken, Christian Ledig, and Zehan Wang. Is the deconvolution layer the same as a convolutional layer? *arXiv preprint arXiv:1609.07009*, 2016.
- [35] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016.
- [36] Min Wang, Baoyuan Liu, and Hassan Foroosh. Factorized convolutional neural networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 545–553, 2017.
- [37] 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 Conference on Computer Vision and Pattern Recognition*, pages 6849–6857, 2019.
- [38] Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE Transactions on Image Processing*, 22(9):3538–3548, 2013.

- [39] Wenjing Wang, Chen Wei, Wenhan Yang, and Jiaying Liu. Gladnet: Low-light enhancement network with global awareness. In *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 751–755. IEEE, 2018.
- [40] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference*, 2018.
- [41] G Westheimer. Spatial interaction in the human retina during scotopic vision. *The Journal of physiology*, 181(4):881–894, 1965.
- [42] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4820–4828, 2016.
- [43] Kai-Fu Yang, Xian-Shi Zhang, and Yong-Jie Li. A biological vision inspired framework for image enhancement in poor visibility conditions. *IEEE Transactions on Image Processing*, 29:1493–1506, 2019.
- [44] Zhenqiang Ying, Ge Li, Yurui Ren, Ronggang Wang, and Wenmin Wang. A new low-light image enhancement algorithm using camera response model. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 3015–3022, 2017.
- [45] Shun-Yuan Yu and Hong Zhu. Low-illumination image enhancement algorithm based on a physical lighting model. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(1):28–37, 2017.
- [46] J. Zhang, R. Liu, L. Ma, W. Zhong, X. Fan, and Z. Luo. Principle-inspired multi-scale aggregation network for extremely low-light image enhancement. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2638–2642, 2020.
- [47] Shuo Zhang, Youfang Lin, and Hao Sheng. Residual networks for light field image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11046–11055, 2019.
- [48] 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*, pages 1632–1640, 2019.
- [49] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018.