Noise Map Guided Inpainting Network for Low-Light Image Enhancement

Zhuolong Jiang[1], Chengzhi Shen[2], Chenghua Li[3(✉)], Hongzhi Liu[4], and Wei Chen[1(✉)]

[1] School of Artificial Intelligence and Computer Science, Jiangnan University, Wuxi, China
wchen_jdsm@jiangnan.edu.cn
[2] Nanjing University of Aeronautics and Astronautics, Nanjing, China
[3] Institute of Automation, Chinese Academy of Sciences, Beijing 100190, China
lichenghua2014@ia.ac.cn
[4] Department of Electronic Engineering, Ocean University of China, Qingdao, China

Abstract. Capturing images in a low-light environment are usually bothered with problems such as serious noise, color degradation, and images underexposure. Most of the low-light image enhancement approaches cannot solve the problem of the loss of the details of the result caused by noise. Inspired by the image inpainting task, we propose a novel Noise-map Guided Inpainting Network (NGI-Net) that introduces inpainting modules to restore missing information. Specifically, the algorithm is divided into two stages. Stage I decomposes input images into a reflection map, an illumination map, and a noise map inspired by the Retinex theory. These maps are passed through Stage II to fine-tune the color and details of the images based on a designed feature enhance group and a selective kernel enhance module. Experiments on real-world and synthesized datasets demonstrate the advantages and robustness of our method. The source code of our method is public in https://github.com/JaChouSSS/NGI-Net.

Keywords: Low-light image enhancement · Inpainting · Retinex decomposition · Selective kernel enhance module

1 Introduction

Capturing images with poor illumination is an issue that people have been trying to solve for a long time, which also widely exists in other tasks such as object detection and video surveillance. Restoring visual information from such images

Z. Jiang—Student.

This work was supported in part by the National Key R&D Program of China under Grant 2017YFB0202303, in part by the National Natural Science Foundation of China under Grand 61602213 and 61772013.

H. Ma et al. (Eds.): PRCV 2021, LNCS 13021, pp. 201–213, 2021.
https://doi.org/10.1007/978-3-030-88010-1_17

is hard due to insufficient detail information, numerous noise, color degradation, low visibility, and low contrast.

For years, researchers have tried to handle the low-light image enhancement task. Traditionally, the histogram-equalization-based methods [1] enhance the images by expanding the dynamic range, but they ignore the distribution of image brightness. Thus, not every region in the images can be effectively improved. Retinex-based approaches [8,9,20] decompose images into reflection maps and illumination maps, but they tend to amplify noise and cause blurry edges. Recently, there are Deep Learning algorithms trained on paired data [15,26,29], unpaired data [7], and those in the zero-reference style [5].

However, the output images of these methods still have problems such as lack of detail, color deviation, and residual noise.

In this paper, we propose a two-stage Noise Map Guided Inpainting Network to enhance low-light images, and at the same time restore image details. To recover detailed information in the extremely dark images, we need to know the distribution of noise as guidance. Specifically, we use a decomposition network as Stage I to decouple the input image, from which we obtain a reflection map with preliminary restoration, as well as a noise map that contains the original noise information. In Stage II, we designed a Feature Enhancement Module (FEM) that extracts valid image features to recover the color and exposure of the reflection map. Then we use the inpainting module to combined the reflection map with the mask generated by the noise map to repair damaged image features.

The main contributions of this paper can be summarized as follows:

1) As far as we know, it is the first time that inpainting tasks are introduced for low-light image enhancement. We also design the inpainting module to recover missing information.
2) We decouple the low-light image enhancement problem into two subtasks and designed a two-stage low-light image enhancement network. Stage I is responsible for decomposing the images, while Stage II further enhances and restores the outputs from Stage I.
3) We design the Feature Enhancement Module (FEM) and the selective kernel enhance module for feature extraction and improvement.

2 Related Works

2.1 Low-Light Image Enhancement

Most of the early low-light image enhancement algorithms use histogram equalization [1] to adjust the global brightness and contrast of the images. Retinex-based algorithms decompose the images into reflection maps and illumination maps. Single-Scale Retinex (SSR) [9] methods adjust the outputs by changing single-scale Gaussian kernels, but they are prone to halo in highlight areas. To mitigate this issue, Multi-Scale Retinex (MSR) [10] methods adopt multi-scale Gaussian kernels. Multiscale Retinex With Color Restoration (MSRSR) [8] tackles color distortions by adding a non-linear restoration factor to MSR to adjust

the numerical ratio of image channels. However, these methods have great limitations, and it is difficult to recover complex scenes and images with strong noise.

With the rapid development of deep learning in recent years, many data-driven low-light image enhancement approaches have been raised. Lore et al. [15] for the first time applied the deep learning method to image enhancement with a deep encoder-decoder structure. Zhang et al. [29] adopted deep neural networks based on Retinex theory to reconstruct the reflection and illumination maps respectively. Due to the scarce of paired image datasets, Jiang et al. [7] proposed an unsupervised network that can be trained on unpaired data. Yang et al. [26] proposed a semi-supervised learning framework and combined it with adversarial learning to improve the perceptual quality of images in poor illumination. Although these methods perform decently in some cases, it is still challenging to solve the problem of information loss and color deviation during recovering extremely dark images.

2.2 Image Inpainting

Traditional image inpainting techniques can be categorified into diffusion-based and patch-based methods. Diffusion-based algorithms [2] propagate inward the image content around the boundary to fill in the missing region. Patch-wise [3] approaches search for the most suitable patch among undamaged regions and fill in the missing one. But both of them cannot generate content with semantic information.

Among DL-based approaches, Pathak et al. [19] proposed an encoder-decoder framework and introduced Generative Adversarial Networks to assist the training process. Liu et al. [14] for the first time presented a Partial Convolution and mask update mechanism for restoring irregular holes, which improved color differences and artifacts. Yu et al. [27] proposed a Gated Convolution method based on [14]. It makes the mask update learnable, strengthens the expressive ability of the network, speeds up the training process, and improves model flexibility. Nazeri et al. [18] integrated the edge generator and the image inpainting network and introduced an image completion algorithm by using edge details as prior knowledge. Inspired by previous works, we design the Inpainting Module for detailed repairs.

3 Method

As illustrated in Fig. 1, our network is divided into two stages: the Decomposition Stage and the Restoration Stage. The Decomposition Stage contains a reflection branch, an illumination branch, and a noise generation branch. Among them, both the reflection branch and the noise generation branch adopt the classic encoding-decoding structure, with the noise generation branch(NB) composed of residual blocks. Through this stage, an input image will be decomposed into

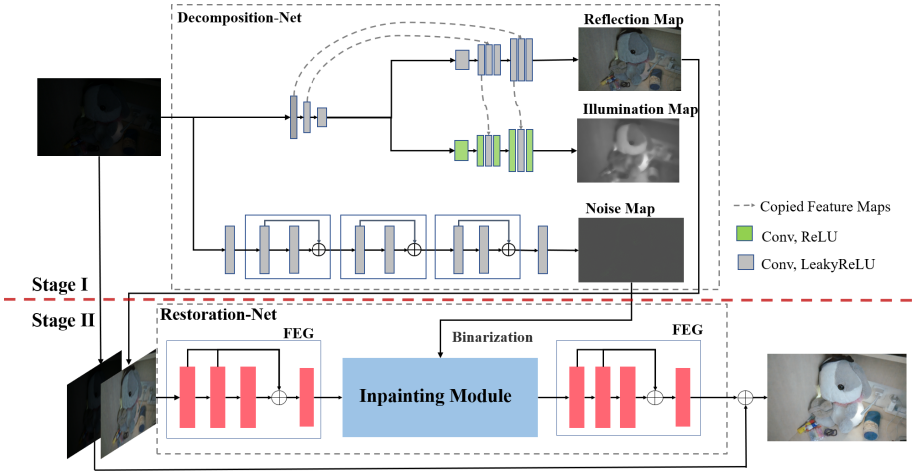

Fig. 1. The pipeline of our NGI-Net model, which contains a Decomposition-Net and a Restoration-Net. The Decomposition-Net contains a reflection branch, an illumination branch, and a noise map generation branch. The Restoration-Net consists of feature enhance groups (FEG) and an inpainting module. The reflection map and noise map from Stage I will be fused with the original image for better restoration.

a reflection map, an illumination map, and a noise map, while masks will be obtained.

The Restoration Stage is made of feature enhance groups and an inpainting module. The FEG is used to make the most of feature information, which can be broken down into sub-modules, named selective kernel enhance module. The inpainting module aims to repair image features with the guidance of first-stage outputs. We will further introduce our network in this section.

3.1 Stage I: Decomposition

In classic Retinex models, an image M is decomposed into the reflection map R and the illumination map I. But they fail in low-light image enhancement because of large amounts of noise. Thus, we adopt the robust Retinex [12] mechanism in the Decomposition Stage, which decomposes images M into reflection map R, illumination map I, and noise map N:

$$M = R \cdot I + N \tag{1}$$

To make more effective use of feature information, the reflection branch (RB) and illumination branch (IB) share the same encoder, which contains 3 convolution layers. A max-pooling layer is added before each convolution to reduce parameters. In the RB, a bilateral upsampling layer with factor 2 is added before each decoder. Skip connections are adopted inside the RB, as well as between the RB and IB with Sigmoid layer. The NB contains 2 convolution layers and 3 Res-blocks, with LeakyReLU as the activation function.

We use the VGG [21] based perceptual loss [11] for calculating $\mathcal{L}_{vgg}$ to constrain the reconstruction error $\mathcal{L}_{rec}$, where M_l and M_h represent the input image and the ground-truth image. They can be decomposed into the reflection map R_l, R_h, the illumination map I_l, I_h, and the noise map N_l, N_h respectively. $\|\cdot\|_1$ is the L1 loss. Further, the illumination map should be piece-wise smooth, thus we adopt $\mathcal{L}_{smooth}$ to minimize the error, where $\bigtriangledown_x$ and $\bigtriangledown_y$ represent first-order derivative operators in the horizontal and vertical directions, and ϵ is a non-negative minimal constant. Additionally, the SSIM loss $\mathcal{L}_{ssim}$ measures the structural similarity between two reflection maps. Since N_h does not exist by default, we use a simple noise loss $\mathcal{L}_{noise}$. Finally, the loss function we use in Stage I is denoted as $\mathcal{L}_I$.

$$\mathcal{L}_I = \mathcal{L}_{rec} + 0.02\mathcal{L}_{smooth} + \mathcal{L}_{ssim} + 20\mathcal{L}_{noise}, \tag{2}$$

where

$$\begin{cases} \mathcal{L}_{ssim} = 1 - SSIM(R_l, R_h), \\ \mathcal{L}_{noise} = \|N_h\|_1, \\ \mathcal{L}_{rec} = \mathcal{L}_{vgg}((R_l \cdot I_l + N_l), M_l) + \mathcal{L}_{vgg}((R_h \cdot I_h + N_h), M_h), \\ \mathcal{L}_{smooth} = \left\| \frac{\bigtriangledown x_l}{\max(\bigtriangledown x_l, \epsilon)} + \frac{\bigtriangledown y_l}{\max(\bigtriangledown y_l, \epsilon)} \right\|_1 + \left\| \frac{\bigtriangledown x_h}{\max(\bigtriangledown x_h, \epsilon)} + \frac{\bigtriangledown y_h}{\max(\bigtriangledown y_h, \epsilon)} \right\|_1. \end{cases} \tag{3}$$

To obtain the mask corresponding to the noise map, we first convert the map into grayscale, and then use the adaptive threshold [17] for binarization.

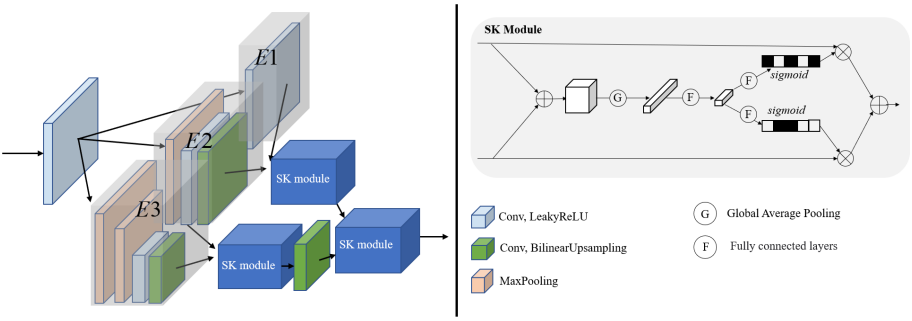

Fig. 2. The structure of the selective kernel enhances modules made of CNN and basic SK modules. SKEM conducts convolution in 3 branches before passing them into SK modules for fusion.

3.2 Stage II: Restoration

Output images from Stage I are bothered by color distortion and detail deficiency, which is analyzed in Sect. 4.3. To handle such issues, we designed the

Restoration Stage that consists of feature enhance groups and an inpainting module and inputs the original input image with a corresponding reflection map. A FEG consists of several sub-modules, named selective kernel enhance module.

Selective Kernel Enhance Module (SKEM). Inspired by previous works on SK [13], we propose the selective kernel enhance modules that function as the basic module to make up the FEG.

While human visual cortex neurons can change their receptive field according to different stimulations [13], CNN simulates this adaptive process by selecting and fusing multi-scale features. Such property is beneficial for the network to expand receptive fields, while in the meantime strengthening feature expression abilities.

As shown in Fig. 2, the input features first go through 1 convolution layer for preliminary extraction, and then respectively go through 3 convolution branches $E1$, $E2$, and $E3$ to obtain features of different scales, where $E1$ is composed of a convolution layer and an activation function, $E2$ adds a max-pooling layer before the convolution layer, and $E3$ adds 2 max-pooling layers to get broader receptive fields. The dimension of output features from SK modules will not be the same. Therefore, we perform channel dimensionality reduction and feature up-sampling through transposed convolution and use the SKEM to fuse them sequentially.

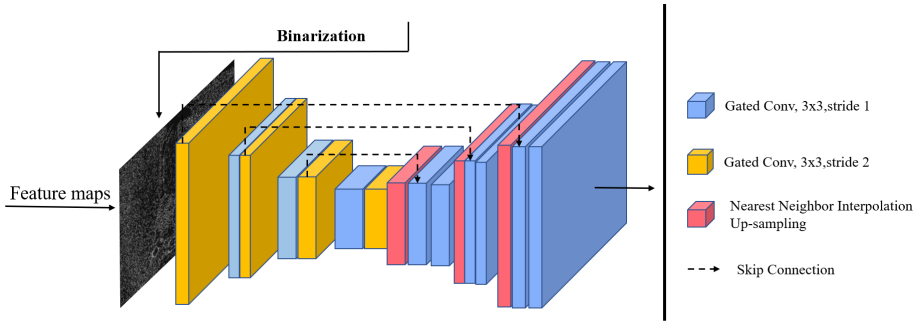

Fig. 3. The structure of our Inpainting Module in the encoder-decoder fashion. We feed the binarized mask and feature maps as input. Gated convolution blocks between the encoder and decoder are connected with skip connections. Nearest Neighbor Interpolation up-sampling layers are placed in the decoder.

Feature Enhance Group (FEG). The FEG is designed as a combination of SKEM. For two reasons, we place FEGs on both the left and right sides of the IM: i) to extract input features from Stage I in support of IM; ii) to further enhance the output of IM. Specifically, the FEG is made of 4 SKEMs, where the first two SKEMs output will be passed to the last one in terms of residual learning. It helps stabilize the training process by merging low-level and high-level semantic information.

Inpainting Module (IM). Traditional convolution takes each pixel as valid, which causes its failure in updating masks with empty regions. Partial convolution takes a step forward by treating input regions as "valid" or "invalid", but only updates the mask heuristically, which limits its flexibility and performance. Gated convolution, however, can learn a dynamic feature selection mechanism for each position of each channel in the feature map. We adopt Gated convolution to update the mask, which can be described as:

$$\begin{aligned} \text{Gating}_{y,x} &= \Sigma\Sigma W_g \cdot I \\ \text{Feature}_{y,x} &= \Sigma\Sigma W_f \cdot I \\ O_{y,x} &= \phi(\text{Feature}_{y,x}) \odot \sigma(\text{Gating}_{y,x}) \end{aligned} \tag{4}$$

where W_g and W_f represent two different kernels for updating masks and computing input features respectively. σ is the sigmoid activation function, and ϕ is the ELU activation function.

The structure of the inpainting module is illustrated in Fig. 3, which is constructed in a encoder-decoder fashion. The encoder is made of 4 Gated convolution blocks. A Gated convolution with kernel-size 3 and stride 2 is adopted in replacement of max-pooling layers, in order to make the most of information from the mask. The decoder is constructed with another 3 different Gated convolution blocks with kernel-size 3 and stride 1. Before each decoder, we use the up-sampling by nearest-neighbor interpolation.

Unlike ordinary masks, each point in the binarized mask is equivalent to a small hole area surrounded with effective boundary information, since the noise map is composed of discrete noise points and extremely small noise blocks. So we add skip connections between encoders and the decoders to pass more information for repairing these holes. The overall loss we use in Stage II is denoted as:

$$\mathcal{L}_{II} = \|\mathcal{F}(M_l, R_l, S) - M_h\|_1 + \mathcal{L}_{vgg}(\mathcal{F}(M_l, R_l, S), M_h) \tag{5}$$

where $\mathcal{F}$ denotes the function of the Stage II network, S represent the noise mask.

We use the same L1 loss as in Stage I to accelerate convergence and enhance model robustness during training. We further adopt the VGG loss that strengthens the model's ability on feature perception and detail restoration.

4 Experiment

We use the two datasets to train and test our network. The Low-Light paired dataset (LOL [25]) contains 500 pairs of 400×600 low/high images from real scenes, of which 485 pairs are used for training and the remaining 15 pairs are used for validation. The synthesized SYN [25] dataset includes 1,000 pairs of 384×384 images, of which 950 pairs are used for training and the remaining 50 pairs are used for validation. We used data augmentation during training on the Stage II network.

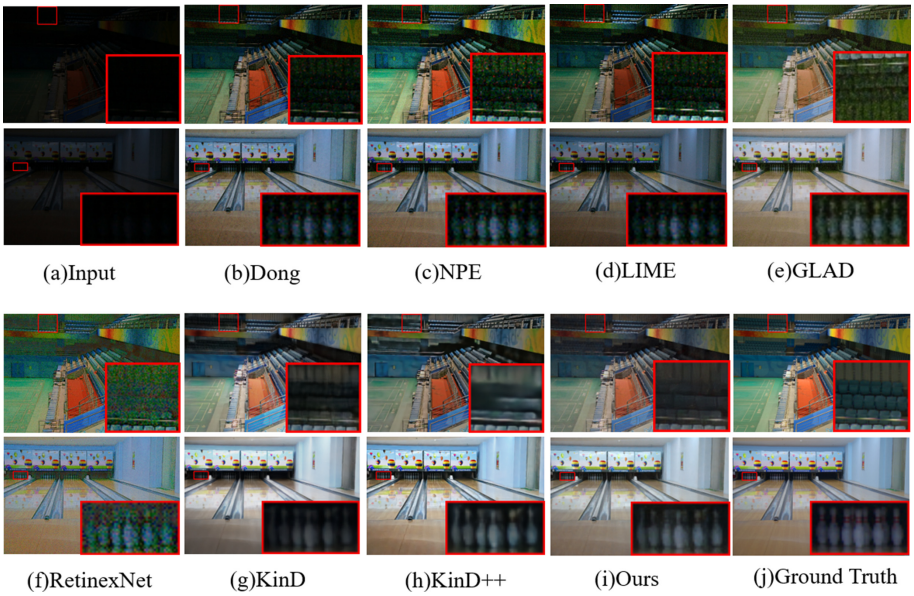

Fig. 4. Qualitative comparison on the output images from the LOL dataset.

4.1 Implementation Details

The training images are randomly cropped and rotated. The batch size is set to 10, the patch-size is set to 96×96. Stages I and II train 300 and 70 epochs respectively. We use the ADAM optimizer, with the learning rate 1e–4 and the cosine annealing learning rate scheduler [16]. Our training takes about 26 h on a single RTX2080Ti GPU.

4.2 Results and Analysis

We evaluate the proposed NGI-Net network and make comparisons with three traditional algorithms: Dong [4], NPE [22], LIME [6], and the state-of-the-art deep-learning methods: GLAD [23], RetinexNet [24], KinD [29], KinD++ [28] in terms of PSNR, SSIM. The quantitative results are shown in Table 1, where ours w/o stage II represents the network with and without stage II and so as to the following two lines. Our low-light enhancement model achieves the best performance among the state-of-the-art methods, with the highest PSNR and SSIM value.

In Fig. 4, we choose a few example images with severe noise and present the restoring results. Obviously, traditional methods fail to remove the noise very well, with some dark areas even remaining unimproved. DL-based methods perform better. RetinexNet has enhanced the extremely dark areas, but with serious noise. Noise also exists in outputs from GLAD, but the algorithm removes some color distortion. Although KinD and KinD++ eliminate some noise, they

Table 1. Experiment results on LOL and SYN dataset.

	LOL Dataset		SYN Dataset	
	PSNR	SSIM	PSNR	SSIM
Dong [4]	16.72	0.4781	16.84	0.7711
NPE [22]	16.97	0.4835	16.47	0.7770
LIME [6]	14.22	0.5203	17.36	0.7868
GLAD [23]	19.72	0.6822	18.05	0.8195
RetinexNet [24]	16.57	0.3989	17.11	0.7617
KinD [29]	20.38	0.8240	18.30	0.8390
KinD++ [28]	21.80	0.8284	19.54	0.8419
Ours w/o stage II	15.68	0.7070	–	–
Ours w/o VGG loss	22.79	0.8079	–	–
Ours w/o inpainting	23.33	0.8319	–	–
Ours	**24.01**	**0.8377**	**26.01**	**0.9366**

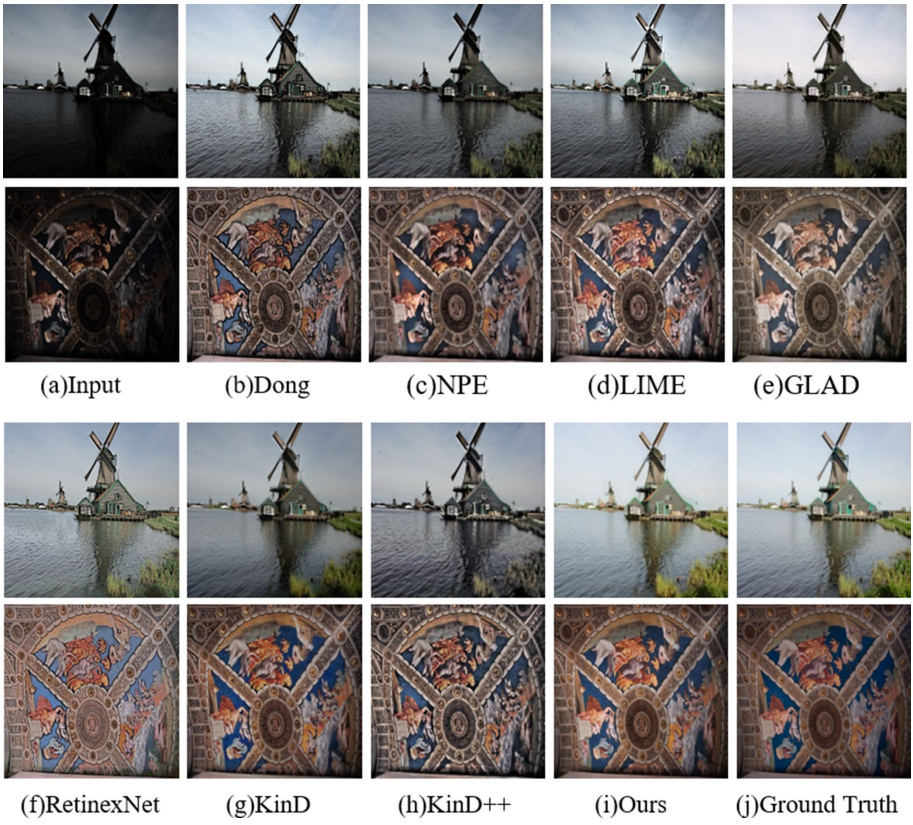

Fig. 5. Qualitative comparison on the output images from the SYN dataset.

have different degrees of blurred boundaries and loss of details. In comparison, our method is more stable that simultaneously removes noise and restores details in extremely dark regions.

Figure 5 mainly illustrates some normal scenarios of low-light images. The results show that traditional algorithms generally have problems such as under-exposure and chromatic aberration. Among deep learning methods, GLAD and RetinexNet generate the most serious color deviation, while KinD and Kind++ perform well on controlling exposure. Our method has good performances on exposure and contrast at the same time, whose outputs are the most similar to ground-truth. Furthermore, Table 1 demonstrate that our network performs better on the LOL and SYN datasets in comparison to others.

4.3 Ablation Study

In this section, we make ablation experiments to evaluate the effectiveness of the proposed components, including the two stage structure, the VGG loss, and the inpainting module. The results of ablation experiments are summarized in Table 1. By only using stage I, we build a baseline model, which achieves 15.68 dB on PSNR and 0.7070 on SSIM. When adding the stage II module, it achieves an increase of 8.33 dB on PSNR and 0.1307 on SSIM, which means the stage II could extract better features to refine the detail information lost in stage I. The effectiveness of VGG loss and the inpainting module can be evaluated by comparing the last three rows in Table 1. It achieves better performance benefitting from the VGG loss and the inpainting process, which means the reconstruction of the detail information is very important to the low-light image enhancement task.

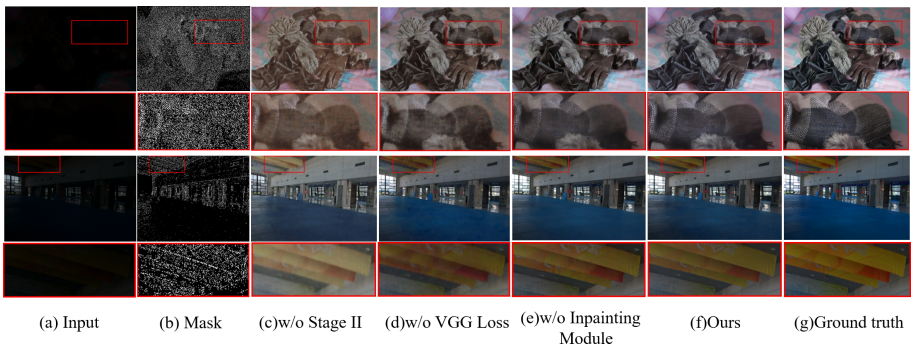

Fig. 6. Visual comparison of the ablation study.

Figure 6 shows the visual changes of the different model structures. For example, Fig. 6(c) shows outputs from a model with/without Stage II that are influenced by severe color deviation and blurry details. After adding Stage II,

color deviations can be significantly restored and high-frequency details can be enhanced. Figure 6(d) describes outputs from a network that is not trained with/without VGG loss. Sample images are of poor color and quality if VGG loss is not used. For example, the edges of the image become blurry and contain amounts of noise. VGG loss helps to restore the details and colors of the image. As mentioned earlier, the inpainting module is used to repair dark regions with information hidden in the noise map. We simply remove inpainting module in Stage II for comparisons. As shown in Fig. 6(e), texture details on the glove drop significantly when the inpainting module is removed, which verifies IM's effectiveness. Moreover, the effectiveness of the feature enhancing group is justified by comparing results from the absence of Stage II and the absence of inpainting module.

5 Conclusion

In this paper, we propose a Noise-map Guided Inpainting Network (NGI-Net) to restore visual information from low-light images. For the first time, we combine the low-light image enhancement task with the inpainting technique, which shows great potential and is more reasonable out of intuition. Our network has two stages. Stage I decomposes input images into the reflection map, illumination map, and noise map. Then Stage II deals with color distortion and details deficiency by feature enhancing and inpainting. The results demonstrate the advantages of our methods based on real-world datasets quantitatively and qualitatively.

However, there are still some limitations to our works. The noise map still cannot separate all the noise regions, which restrains the network's capability on detail restoration. In future works, we will focus on enhancing the semantic information of dark regions of the image to design robust approaches for more sophisticated low-light image enhancement tasks.

References

1. Arici, T., Dikbas, S., Altunbasak, Y.: A histogram modification framework and its application for image contrast enhancement. IEEE Trans. Image Process. **18**, 1921–1935 (2009)
2. Ballester, C., Bertalmío, M., Caselles, V., Sapiro, G., Verdera, J.: Filling-in by joint interpolation of vector fields and gray levels. IEEE Trans. Image Process. Publication IEEE Signal Process. Soc. **10**(8), 1200–11 (2001)
3. Criminisi, A., Pérez, P., Toyama, K.: Region filling and object removal by exemplar-based image inpainting. IEEE Trans. Image Process. **13**, 1200–1212 (2004)
4. Dong, X., Pang, Y., Wen, J.: Fast efficient algorithm for enhancement of low lighting video. 2011 IEEE International Conference on Multimedia and Expo, pp. 1–6 (2011)
5. Guo, C., et al.: Zero-reference deep curve estimation for low-light image enhancement. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1777–1786 (2020)

6. Guo, X., Li, Y., Ling, H.: Lime: low-light image enhancement via illumination map estimation. IEEE Trans. Image Process. **26**, 982–993 (2017)
7. Jiang, Y., et al.: Enlightengan: deep light enhancement without paired supervision. IEEE Trans. Image Process. **30**, 2340–2349 (2021)
8. Jobson, D.J., Rahman, Z., Woodell, G.A.: A multiscale retinex for bridging the gap between color images and the human observation of scenes. IEEE Trans. Image Process. Publication IEEE Signal Process. Soc. **6**(7), 965–76 (1997)
9. Jobson, D.J., Rahman, Z., Woodell, G.A.: Properties and performance of a center/surround retinex. IEEE Trans. Image Process. Publication IEEE Signal Process. Soc. **6**(3), 451–62 (1997)
10. Jobson, D.J., Rahman, Z.u., Woodell, G.A.: A multiscale retinex for bridging the gap between color images and the human observation of scenes. IEEE Trans. Image Process. **6**(7), 965–976 (1997)
11. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. ArXiv abs/1603.08155 (2016)
12. Li, M., Liu, J., Yang, W., Sun, X., Guo, Z.: Structure-revealing low-light image enhancement via robust retinex model. IEEE Trans. Image Process. **27**, 2828–2841 (2018)
13. Li, X., Wang, W., Hu, X., Yang, J.: Selective kernel networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 510–519 (2019)
14. Liu, G., Reda, F., Shih, K.J., Wang, T., Tao, A., Catanzaro, B.: Image inpainting for irregular holes using partial convolutions. ArXiv abs/1804.07723 (2018)
15. Lore, K.G., Akintayo, A., Sarkar, S.: Llnet: a deep autoencoder approach to natural low-light image enhancement. Pattern Recogn. **61**, 650–662 (2017)
16. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. arXiv: Learning (2017)
17. McAndrew, A.: An introduction to digital image processing with matlab notes for scm2511 image processing. School of Computer Science and Mathematics, Victoria University of Technology **264**(1), 1–264 (2004)
18. Nazeri, K., Ng, E., Joseph, T., Qureshi, F., Ebrahimi, M.: Edgeconnect: Generative image inpainting with adversarial edge learning. ArXiv abs/1901.00212 (2019)
19. Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., Efros, A.A.: Context encoders: feature learning by inpainting. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2536–2544 (2016)
20. Rahman, Z., Jobson, D.J., Woodell, G.A.: Multi-scale retinex for color image enhancement. In: Proceedings of 3rd IEEE International Conference on Image Processing 3, vol. 3, pp. 1003–1006 (1996)
21. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR abs/1409.1556 (2015)
22. Wang, S., Zheng, J., Hu, H., Li, B.: Naturalness preserved enhancement algorithm for non-uniform illumination images. IEEE Trans. Image Process. **22**, 3538–3548 (2013)
23. Wang, W., Wei, C., Yang, W., Liu, J.: Gladnet: low-light enhancement network with global awareness. In: 2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018), pp. 751–755 (2018)
24. Wei, C., Wang, W., Yang, W., Liu, J.: Deep retinex decomposition for low-light enhancement. In: BMVC (2018)
25. Wei, C., Wang, W., Yang, W., Liu, J.: Deep retinex decomposition for low-light enhancement. arXiv preprint arXiv:1808.04560 (2018)

26. Yang, W., Wang, S., Fang, Y., Wang, Y., Liu, J.: From fidelity to perceptual quality: a semi-supervised approach for low-light image enhancement. 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 3060–3069 (2020)
27. Yu, J., Lin, Z.L., Yang, J., Shen, X., Lu, X., Huang, T.: Free-form image inpainting with gated convolution. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 4470–4479 (2019)
28. Zhang, Y., Guo, X., Ma, J., Liu, W., Zhang, J.: Beyond brightening low-light images. International Journal of Computer Vision, pp. 1–25 (2021)
29. Zhang, Y., Zhang, J., Guo, X.: Kindling the darkness: a practical low-light image enhancer. In: Proceedings of the 27th ACM International Conference on Multimedia, pp. 1632–1640 (2019)