DarkIR: Robust Low-Light Image Restoration
— Supplementary Material —

Daniel Feijoo [1†], Juan C. Benito [1], Alvaro Garcia [1], Marcos V. Conde [1,2†]
[1] Cidaut AI, Valladolid, Spain
[2] Computer Vision Lab, University of Würzburg, Germany
https://github.com/cidautai/DarkIR † Corresponding Authors {danfei,marcos.conde@cidaut.es}

1. Additional Implementation Details

Our implementation is based on PyTorch [12]. We train DarkIR (following LEDNet [19]) on the LOLBlur dataset. During training, we randomly crop 384×384 patches, and apply standard flip and rotation augmentations. The mini-batch size is set to 32 using an H100 GPU.

As our optimizer we use AdamW [8] by setting $\beta_1 = 0.9$, $\beta_2 = 0.9$ and weight decay to $1e^{-3}$. The learning rate is initialized to $5e^{-4}$ and is updated by the cosine annealing strategy [7] to a minimum of $1e^{-6}$. We repeat this configuration for re-training the other methods in LOLBlur dataset. Note that we use the official open-source implementation of the other methods, or previously reported results.

The **multi-task** model was trained using the same setup. The only difference is the use of LOLv2 and LSRW as additional datasets. This model achieves essentially state-of-the-art results on real low-light enhancement benchmarks, while maintaining the performance on LOLBlur.

2. Additional ablation Studies

We studied the influence of the optimization losses. The results can be seen on Table A, where we can check that by introducing the L_{edge}, L_{lol} and L_{percep} the model achieves the best combination of distortion and perceptual metrics.

Table A. Ablation study on our loss functions. We train DarkIR using different loss setups. Adding the perceptual loss (L_{percep}), edge loss (L_{edge}) and the architecture guiding loss (L_{lol}) helps to improve the overall performance.

	PSNR↑	SSIM↑	LPIPS↓
L_{pixel}	26.34	0.856	0.205
L_{pixel} + L_{lol}	26.19	0.861	0.197
L_{pixel} + L_{lol} + L_{edge}	<u>26.717</u>	0.874	0.182
L_{pixel} + L_{percep} + L_{edge}	26.61	**0.877**	**0.171**
L_{pixel} + L_{lol} + L_{edge} + L_{percep}	**26.9**	<u>0.874</u>	<u>0.176</u>

Table B. Ablation study on the feature propagation between encoder and decoder. We found simple addition to be optimal.

	Params↓ (M)	MACs↓ (G)	PSNR↑	SSIM↑	LPIPS↓
CurveNLU	4.05	14.14	26.55	0.872	0.176
CurveNLU-DepthWise	3.33	7.42	26.64	0.872	0.177
1DLUT	3.39	7.87	26.69	0.873	**0.175**
1DLUT-double	3.49	8.9	26.63	**0.875**	<u>0.175</u>
Single Addition	**3.31**	**7.25**	**26.9**	<u>0.874</u>	0.176

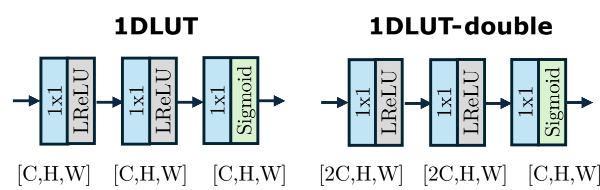

Figure A. Neural blocks proposed for the feature propagation. The difference between both is the presence of a channel expansion in **1DLUT-double**.

In addition, we studied different skip connections for the feature propagation between encoder and decoder. The proposed feature propagation takes the form of:

$$y = f_{prop}(enc_{feat}) + dec_{feat} \tag{1}$$

where f_{prop} is the proposed feature block applied to the encoder features (enc_{feat}) added to the decoder features (dec_{feat}). Besides the CurveNLU proposed by LEDNet [19] we evaluate the results obtained by using only depth-wise convolutions in this given block. To sum up we incorporate a variation of this block that uses only point-wise convolutions, resembling the behaviour of a look-up table (LUT). Figure A represents the proposed 1DLUT variations. In Table B the results of this ablation study are showcased. We see that the single addition, i.e $f_{prop} = Identity$ gets the best performance, so we did not consider adding any of the discussed blocks to the DarkIR architecture.

Table C. Quantitative comparison on five **real-workd unpaired LLIE** datasets using the perceptual quality metrics BRISQUE [10] and NIQE [11]. We use reference results from [17].

LLIE	**DICM**		**LIME**		**MEF**		**NPE**		**VV**	
Unpaired	BRISQUE↓	NIQE↓	BRISQUE↓	NIQE↓	BRISQUE↓	NIQE↓	BRISQUE↓	NIQE↓	BRISQUE↓	NIQE↓
KinD [18]	48.72	5.15	39.91	5.03	49.94	5.47	36.85	4.98	50.56	4.30
ZeroDCE [2]	27.56	4.58	<u>20.44</u>	5.82	17.32	4.93	20.72	4.53	34.66	4.81
RUAS [6]	38.75	5.21	27.59	4.26	23.68	3.83	47.85	5.53	38.37	4.29
SNR-Net [16]	37.35	4.71	39.22	5.74	31.28	4.18	26.65	4.32	78.72	9
CIDNet [17]	21.47	3.79	**16.25**	<u>4.13</u>	**13.77**	<u>3.56</u>	<u>18.92</u>	**3.74**	<u>30.63</u>	**3.21**
DarkIR-mt	**18.69**	**3.76**	21.62	**4.07**	<u>13.90</u>	**3.45**	**12.88**	<u>3.99</u>	**26.87**	<u>3.74</u>

3. More quantitative results in unpaired data

In addition to the results of unpaired Real-LOLBlur datasets, in Table C we present the results of our model in 5 well-known unpaired datasets: LIME [3], DICM [5], MEF [9], NPE [15] and VV [14]. We use the multi-task model trained in LOLBLur and in the LOL datasets. We report BRISQUE [10] and NIQE [11] metrics.

4. More qualitative results in LLIE

As we indicated in the paper, we present more qualitative results in Figures B, C and D. These results showcase the power of our <u>multi-task model</u> for LLIE restoration. Then, in Figures E, F and G we show more qualitative results on the LLIE-Deblurring task.

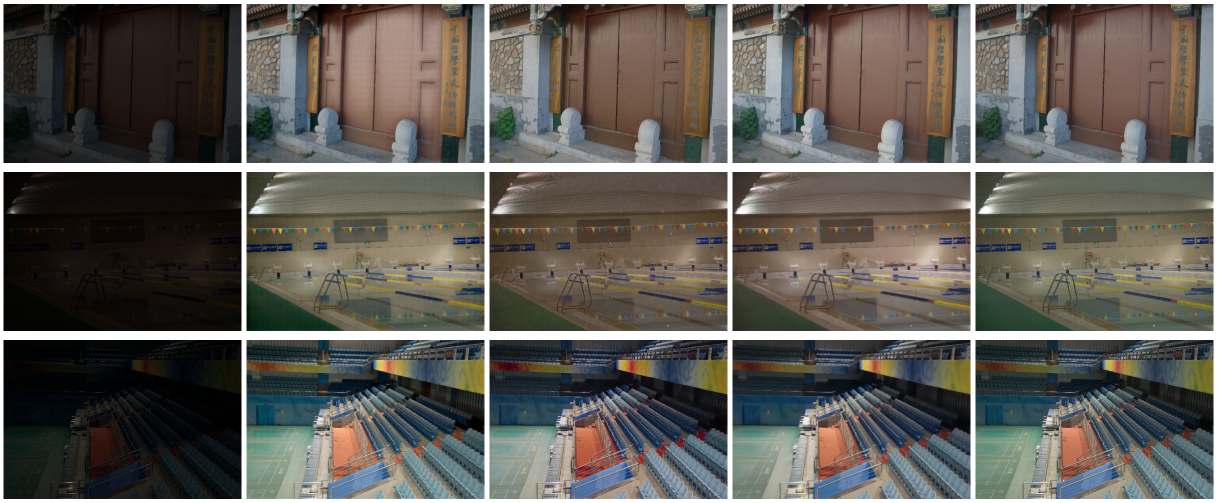

Input SNR-Net RetinexFormer **DarkIR-m** Ground Truth

Figure B. Qualitative results compared with state of the art method RetinexFormer [1] and SNR-Net [16] on **LOLv2-Real**.

Input SNR-Net RetinexFormer **DarkIR-m** Ground Truth

Figure C. Qualitative results compared with state of the art method RetinexFormer [1] and SNR-Net [16] on **LOLv2-Synthetic**.

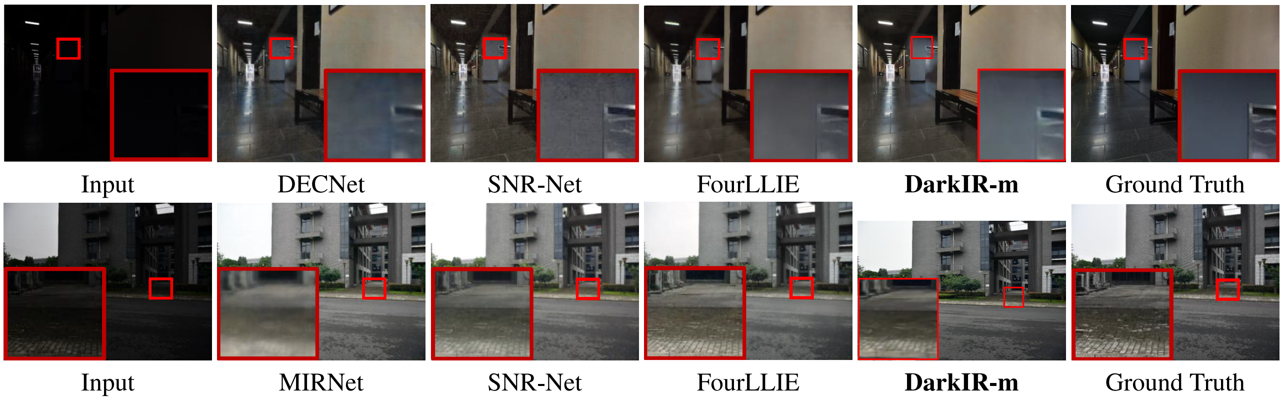

Figure D. Qualitative results on the real-world dataset **LSRW-Huawei** [4] (top row) and **LSRW-Nikon** [4] (bottom row).

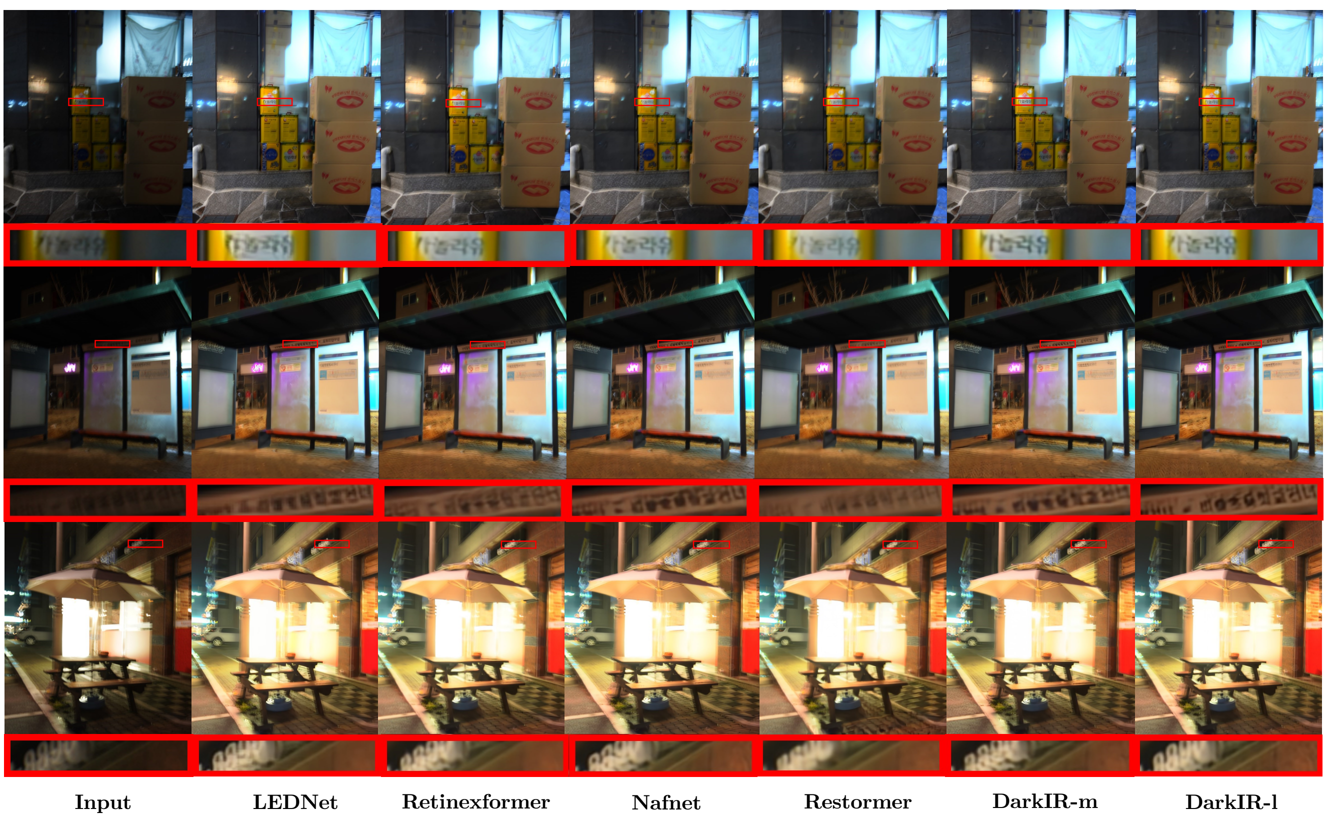

Figure E. Qualitative results on **RealBlur-Night** [13] images. (Zoom in for best view).

Figure F. Qualitative results in **Real-LOLBlur** [19] dataset. (Zoom in for best view).

Figure G. Qualitative results in **LOLBlur** [19] dataset. As we can see, **DarkIR** gets sharper and brighter results than the other methods. (Zoom in for best view).

References

[1] Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12504–12513, 2023. 3

[2] Chunle Guo 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 conference on computer vision and pattern recognition (CVPR)*, pages 1780–1789, 2020. 2

[3] 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. 2

[4] Jiang Hai, Zhu Xuan, Ren Yang, Yutong Hao, Fengzhu Zou, Fang Lin, and Songchen Han. R2rnet: Low-light image enhancement via real-low to real-normal network. *Journal of Visual Communication and Image Representation*, 90: 103712, 2023. 4

[5] Chulwoo Lee, Chul Lee, and Chang-Su Kim. Contrast enhancement based on layered difference representation of 2d histograms. *IEEE transactions on image processing*, 22(12): 5372–5384, 2013. 2

[6] Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10561–10570, 2021. 2

[7] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 1

[8] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. 1

[9] 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. 2

[10] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing*, 21(12): 4695–4708, 2012. 2

[11] Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a "completely blind" image quality analyzer. *IEEE Signal processing letters*, 20(3):209–212, 2012. 2

[12] 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. *Advances in neural information processing systems*, 32, 2019. 1

[13] Jaesung Rim, Haeyun Lee, Jucheol Won, and Sunghyun Cho. Real-world blur dataset for learning and benchmarking deblurring algorithms. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020. 4

[14] Vassilios Vonikakis, Rigas Kouskouridas, and Antonios Gasteratos. On the evaluation of illumination compensation algorithms. *Multimedia Tools and Applications*, 77:9211–9231, 2018. 2

[15] 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. 2

[16] Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *CVPR*, 2022. 2, 3

[17] Qingsen Yan, Yixu Feng, Cheng Zhang, Pei Wang, Peng Wu, Wei Dong, Jinqiu Sun, and Yanning Zhang. You only need one color space: An efficient network for low-light image enhancement. *arXiv preprint arXiv:2402.05809*, 2024. 2

[18] Yonghua Zhang, Xiaojie Guo, Jiayi Ma, Wei Liu, and Jiawan Zhang. Beyond brightening low-light images. *International Journal of Computer Vision*, 129:1013–1037, 2021. 2

[19] Shangchen Zhou, Chongyi Li, and Chen Change Loy. Lednet: Joint low-light enhancement and deblurring in the dark. In *ECCV*, 2022. 1, 5, 6