Physically Inspired Dense Fusion Networks for Relighting

Runner Up of NTIRE 2021 Depth-Guided Relighting Challenge

Amirsaeed Yazdani (amiryazdani@psu.edu)
Tiantong Guo (tiantong@ieee.org)
Vishal Monga (vmonga@engr.psu.edu)

Information Processing & Algorithm Laboratory
Pennsylvania State University

June 19 2021

Fusion Strategy

- We propose a fusion strategy for generating the relit image:

Fusion Strategy

- We propose a fusion strategy for generating the relit image:
  - Intrinsic image decomposition:

$$I_{\text{intrinsic-relit}} = \hat{A} \odot \hat{S}$$

Fusion Strategy

- We propose a fusion strategy for generating the relit image:
  - Intrinsic image decomposition:

$$I_{\text{intrinsic-relit}} = \hat{A} \odot \hat{S}$$

  - Direct relighting:

$$f(I) = I_{\text{direct-relit}}$$

Fusion Strategy

- We propose a fusion strategy for generating the relit image:
  - Intrinsic image decomposition:

$$I_{\text{intrinsic-relit}} = \hat{A} \odot \hat{S}$$

  - Direct relighting:

$$f(I) = I_{\text{direct-relit}}$$

- The fused output:

$$\hat{I}_{\textit{relit}} = wI_{\textit{direct-relit}} + (1-w)I_{\textit{intrinsic-relit}}$$

One-to-one Intrinsic Decomposition Direct RelightNet [1]

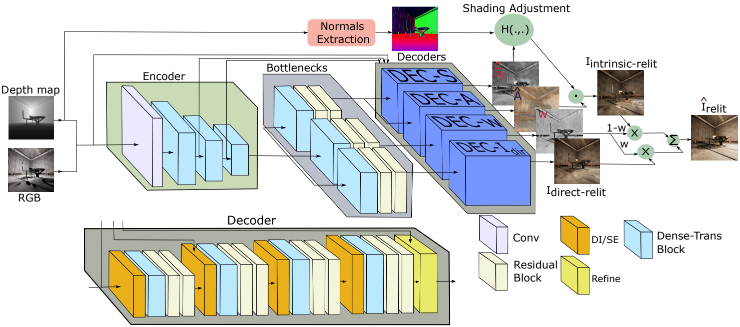

FIGURE – Our proposed OIDDR-Net. The shading estimate is adjusted using the normals of the surfaces in the scene provided by the normal extraction module.

1. For the implementation details please visit: github/Relighting

Any-to-any Multiscale Intrinsic-Direct RelightNet

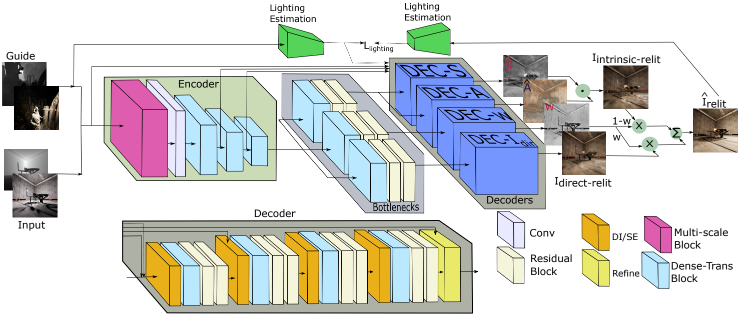

FIGURE – Our proposed AMIDR-Net. The multi-scale block helps the model systematically learn to upsample and downsample the inputs according to the features it needs.

Training Loss

$$\mathcal{L} = \mathcal{L}_{\textit{total}} + \lambda_1 \mathcal{L}_{IID} + \lambda_2 \mathcal{L}_{\textit{direct}} + \lambda_3 \mathcal{L}_{\textit{SSIM}} + \lambda_4 \mathcal{L}_{\textit{lighting}}$$

Training Loss

$$\mathcal{L} = \mathcal{L}_{\textit{total}} + \lambda_1 \mathcal{L}_{IID} + \lambda_2 \mathcal{L}_{\textit{direct}} + \lambda_3 \mathcal{L}_{\textit{SSIM}} + \lambda_4 \mathcal{L}_{\textit{lighting}}$$

$$\mathcal{L}_{\textit{total}} = ||\hat{I}_{\textit{relit}} - Y_{\textit{relit}}||_2^2$$

$$\mathcal{L}_{IID} = ||\hat{A} \odot \hat{S} - Y_{\textit{relit}}||_2^2 + ||\hat{A} - A||_2^2 + ||\hat{S} - S||_2^2$$

$$\mathcal{L}_{\textit{direct}} = ||I_{\textit{direct-relit}} - Y_{\textit{relit}}||_2^2$$

$$\mathcal{L}_{SSIM} = 1 - SSIM(\hat{I}_{\textit{relit}}, Y_{\textit{relit}})$$

$$\mathcal{L}_{\textit{lighting}} = ||g(\hat{I}_{\textit{relit}}) - g(Y_{\textit{relit}})||_2^2 - \sum_{i=1}^{8} Y^i_{\textit{dir-guide}} log(\hat{Y}^i_{\textit{dir}}) - \sum_{j=1}^{5} Y^j_{\textit{color-guide}} log(\hat{Y}^j_{\textit{color}})$$

One-to-one Relighting

Team	Author	**MPS** ↑	SSIM ↑	LPIPS ↓	PSNR ↑	Run-time	Platform	GPU
AICSNTU-MBNet	HaoqiangYang	0.7663	0.6931	0.1605	19.1469	2.88s	PyTorch	Tesla V100
iPAL-RelightNet	auy200	**0.7620**	**0.6874**	**0.1634**	**18.8358**	0.53s	PyTorch	Titan XP
NTUAICS-ADNet	aics	0.7601	0.6799	0.1597	18.8639	2.76s	PyTorch	Tesla V100
VUE	lifu	0.7600	0.6903	0.1702	19.8645	0.23s	PyTorch	P40
NTUAICS-VGG	jimmy3505090	0.7551	0.6772	0.1670	18.2766	2.12s	PyTorch	Tesla V100
DeepBlueAI	DeepBlueAI	0.7494	0.6879	0.1891	19.8784	0.17s	PyTorch	Tesla V100
usuitakumi	usuitakumi	0.7229	0.6260	0.1801	16.8249	0.04s	PyTorch	Tesla V100
MCG-NKU	NK_ZZL	0.7147	0.6191	0.1896	19.0856	0.33s	PyTorch	RTX TITAN
alphaRelighting	lchia	0.7101	0.6084	0.1882	15.8591	0.04s	PyTorch	Tesla K80
Wit-AI-lab	MDSWYZ	0.6966	0.6113	0.2181	17.5740	0.9s	PyTorch	RTX 2080Ti
Couger AI	Sabarinathan	0.6475	0.5469	0.2518	18.2938	0.015s	Tensorflow	GTX 1070

TABLE – NTIRE 2021 Depth-Guided Image Relighting Challenge Track 1 (One-to-one relighting) results. The MPS is used to determine the final ranking.

Any-to-any Relighting

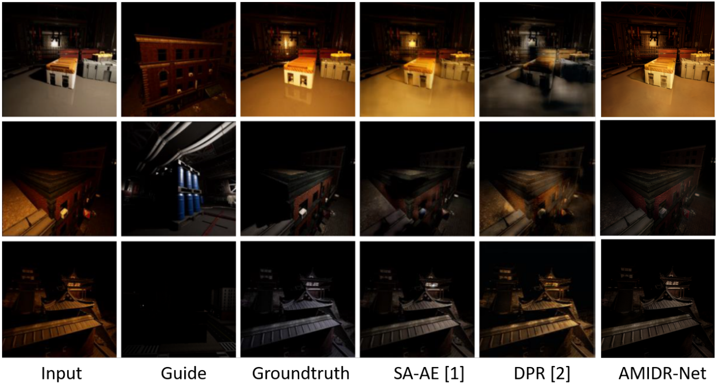

FIGURE – Comparing our AMIDR-Net with other methods on samples from VIDIT'20 [3] dataset.

References

1) H. Hu, X. Huang, Y. Li, and Q. Wang. "SA-AE for any-to-any relighting." Computer Vision – ECCV 2020 Workshops, pages 535–549, Cham, 2020. Springer International Publishing.
2) H. Zhou, S. Hadap, K. Sunkavalli, and D. Jacobs. "Deep single-image portrait relighting." In 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pages 7193–7201, 2019.
3) M. Helou, R. Zhou, J. Barthas, and Susstrunk. "VIDIT: Virtual image dataset for illumination transfer." arXiv preprint arXiv:2005.05460, 2020.

Thank You !

Physically Inspired Dense Fusion Networks for Relighting

PennState College of Engineering
Information Processing & Algorithms Laboratory
Everything can be optimized

http://signal.ee.psu.edu