CURL: Neural Curve Layers for Global Image Enhancement

Sean Moran[1†], Steven McDonagh[2], Gregory Slabaugh[3†]

JP Morgan Chase & Co.[1] Huawei Noah's Ark Lab[2] Queen Mary University of London[3]

25th INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION
Milan, Italy 10 | 15 January 2021

Introduction

Problem

- Final image quality is of pivotal importance in end-user imaging systems such as DSLR and smartphone cameras
- Post-capture enhancement is an ill-posed problem that demands appropriate transformation of properties such as colour, saturation and luminance
- Successful manual enhancement often requires digital artistic **skill, time** and **professional software**

Observations

- Digital artists modify image properties through manual control of *adjustment curves*
- Popular parameterised manual enhancement tools allow for smooth global image adjustments
- Effective image enhancement requires both *local* and *global* adjustment

Research questions

1. Can we automatically estimate, and apply, image *adjustment curves* to improve perceptual quality?
2. Which curves and colour spaces should be considered?
3. Does adjustment application ordering matter?

Poor exposure input | Our approach | Professional digital artist (GT)

Proposed approach and contributions

- CURL: **multi-colour-space neural retouching block**. We learn piece-wise linear scaling curves towards adjusting image properties in a human-interpretable manner
- A **multi-colour-space loss function**: guides sequential and differentiable image transforms in multiple colour spaces (HSV, Lab, RGB)
- TED: **T**ransformed **E**ncoder-**D**ecoder backbone. We modify network backbone architectures by streamlining the use of skip connections towards improving decoder performance
- State-of-the-art performance on three competitive benchmarks

TED and CURL image enhancement: Method Overview

- Our CURL retouching block receives input image features from an arbitrary encoder-decoder backbone network
- TED provides an example multi-scale encoder-decoder backbone, and this pairing defines our system (Method Overview)
- Local pixel adjustments are enabled by the backbone (TED) and the retouching block (CURL) learns *curve layers* that allow global adjustment of image properties

† work carried out at Huawei Noah's Ark Lab

Method Overview

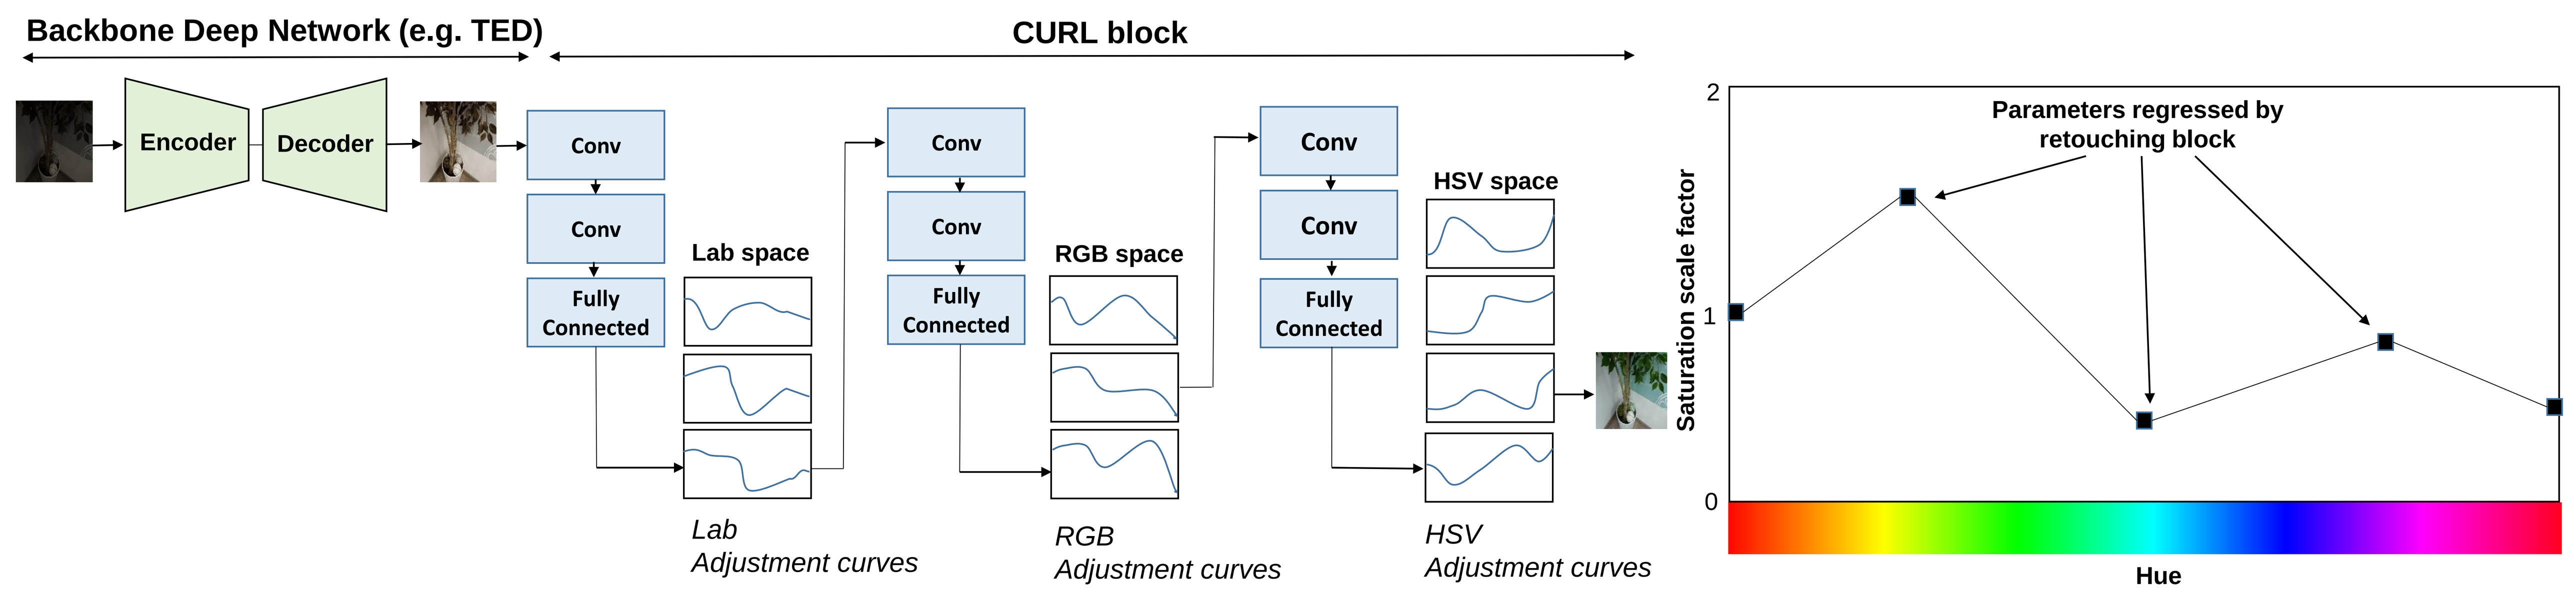

Transformed Encoder/ Decoder (TED) backbone

- Our backbone architecture TED enables local pixel adjustment with a U-Net like structure
- We find global and mid-level contextual consideration to be important; lessening spatial inconsistencies
- Context is accounted for by fusing global and mid-level features using a *multi-scale contextual awareness* (MSCA) connection
- Our *single* skip connection trades-off parameter complexity and image quality
- We highlight that *endowing a skip connection with multi-scale processing* reduces relative parameter counts yet can also improve image quality
- Fusing multiple levels of image context delivers more contextually relevant features for the decoder path

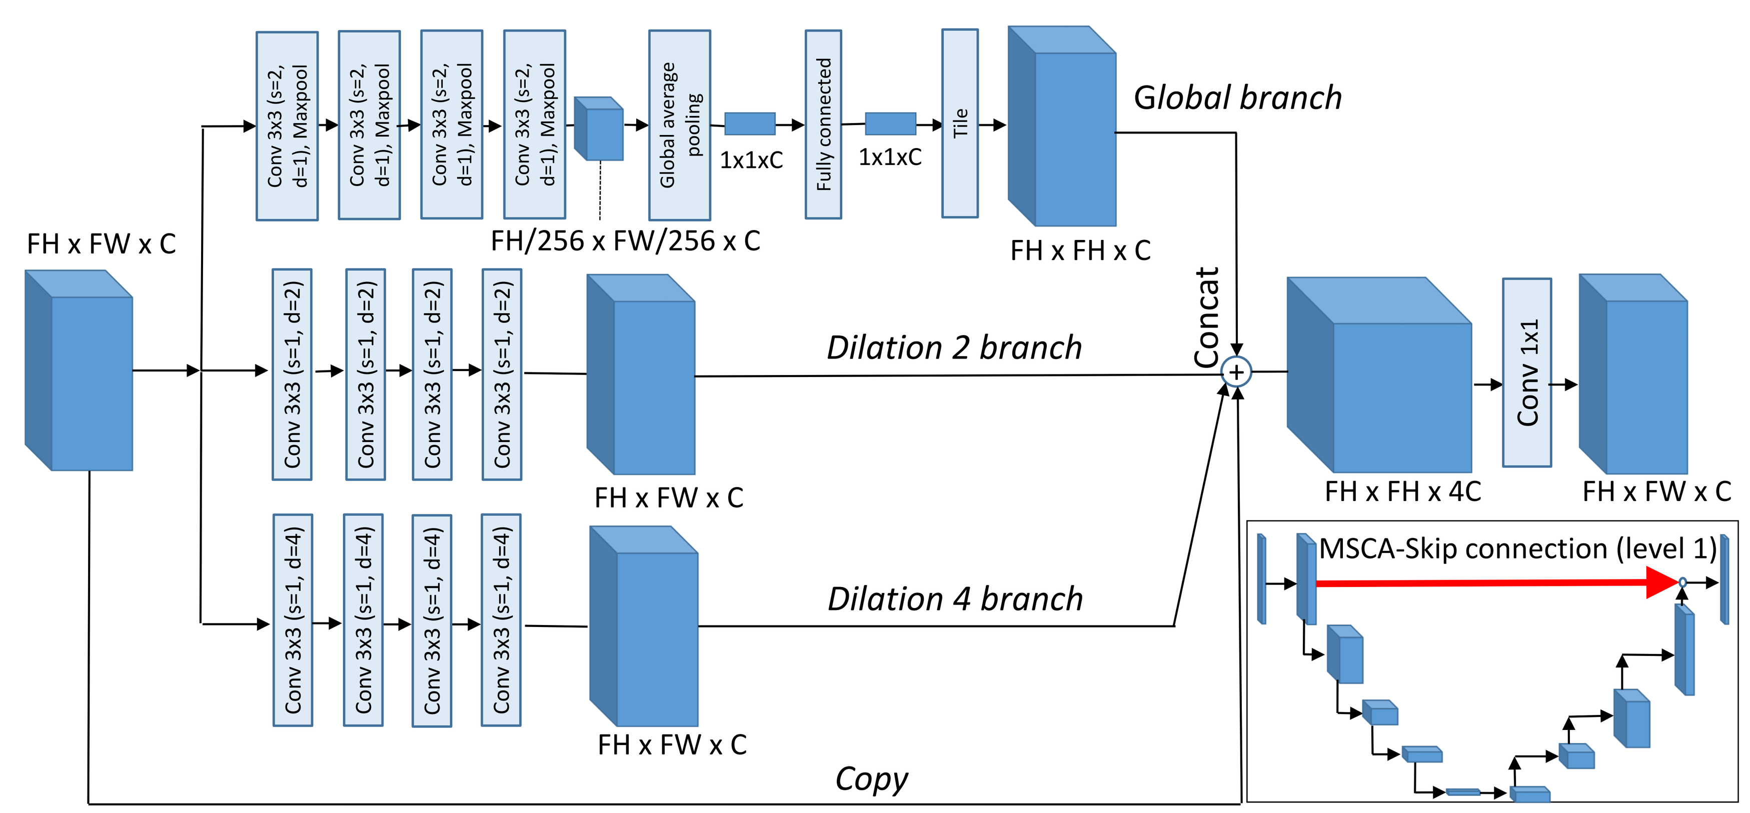

CURL: Neural Curve Layers for Global Image Adjustment

- CURL: our global image colour, saturation and luminance retouching block
- The previous TED network outputs a $H\times W\times C$ tensor, that is fed to the CURL block
- $H\times W\times C$ tensor represents the image to be globally enhanced, concatenated with additional feature maps serving as *neural curve layer* input features
- We perform a regression to acquire curve knot points of piecewise-linear adjustment curves, for multiple colour spaces
- Learned adjustment curves enhance image I by scaling pixels in three colour spaces (HSV, CIELab, RGB)
- CURL therefore regresses expressive curves, used to *scale* rather than *remap* colours *c.f.* existing approaches

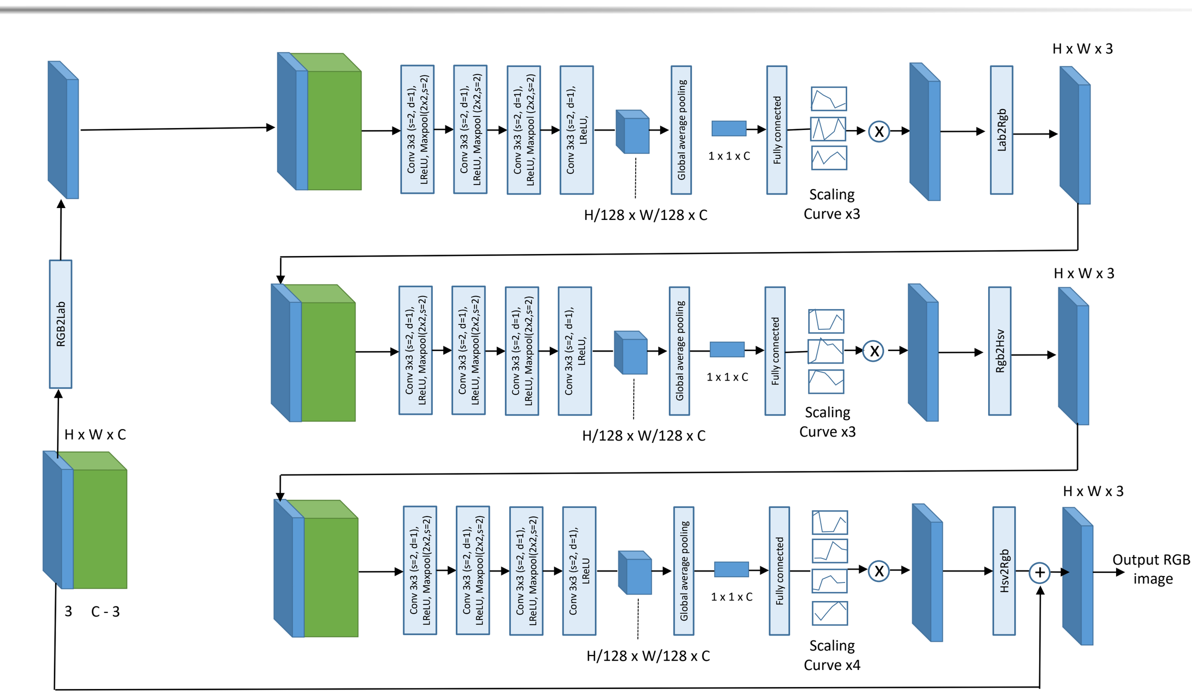

Quantitative results

- **Example experiment**: prediction quality for photographer C enhancement retouchings (groundtruth). **MIT-Adobe 5K** [3] test images
- State-of-the-art performance; PSNR, SSIM, LPIPS
- **Efficient** architecture (neural weight parameter count)

Architecture	PSNR↑	SSIM↑	LPIPS↓	# Params
TED+CURL	**24.04**	**0.900**	**0.583**	1.4M
DPED [4]	21.76	0.871	–	–
8RESBLK [5]	23.42	0.875	–	–
FCN [2]	20.66	0.849	–	–
CRN [1]	22.38	0.877	–	–
U-Net [6]	22.13	0.879	–	–
DPE [3]	23.80	0.900	0.587	3.3M

↓ Please see our paper for additional quantitative results ↓

Qualitative results

DeepISP [7] (28.19 dB) | **TED+CURL (29.37 dB)** | Groundtruth

Example: Samsung S7 *Medium Exposure* dataset

DeepUPE [8] (16.85 dB) | **TED+CURL (23.55 dB)** | Groundtruth

Example: *MIT-Adobe-UPE* dataset

Summary

- We introduce CURL: CURve Layers for learnable image enhancement
- Inspired by digital artists, we learn image retouching adjustment curves
- Exploitation of image representation in three different colours spaces (CIELab, HSV, RGB)
- Extensive experimentation reports state-of-the-art quantitative, qualitative results across a suite of benchmarks

References

[1] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *ICCV*, pages 1511–1520, 2017.

[2] Qifeng Chen, Jia Xu, and Vladlen Koltun. Fast image processing with fully-convolutional networks. In *ICCV*, pages 2497–2506, 2017.

[3] Yu-Sheng Chen, Yu-Ching Wang, Man-Hsin Kao, and Yung-Yu Chuang. Deep photo enhancer: Unpaired learning for image enhancement from photographs with gans. In *CVPR*, June 2018.

[4] Andrey Ignatov, Nikolay Kobyshev, Radu Timofte, Kenneth Vanhoey, and Luc Van Gool. Dslr-quality photos on mobile devices with deep convolutional networks. In *ICCV*, 2017.

[5] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *NeurIPS*, pages 700–708, 2017.

[6] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015.

[7] E. Schwartz, R. Giryes, and A. M. Bronstein. DeepISP: Towards Learning an End-to-End Image Processing Pipeline. *IEEE Transactions on Image Processing*, 28(2):912–923, 2019.

[8] Ruixing Wang, Qing Zhang, Chi-Wing Fu, Xiaoyong Shen, Wei-Shi Zheng, and Jiaya Jia. Underexposed photo enhancement using deep illumination estimation. In *CVPR*, 2019.

Links

Contact: sean.j.moran@gmail.com

Paper: Code: