Colorizing Color Images

Ligeng Zhu and Brian Funt

Simon Fraser University

Introduction

Goal

Remove spatially varying color casts from images

Intuition

(1) Deep learning methods colorize luminance (greyscale) images quite believably

(2) Colorization methods must be encoding knowledge of the world (i.e., sky is blue)

(3) Colorization works for luminance, so why not add color channels too?

(4) Hypothesize that encoded world knowledge will help remove unnatural color casts.

Colorization of Luminance Image Examples [1]

Network Architecture

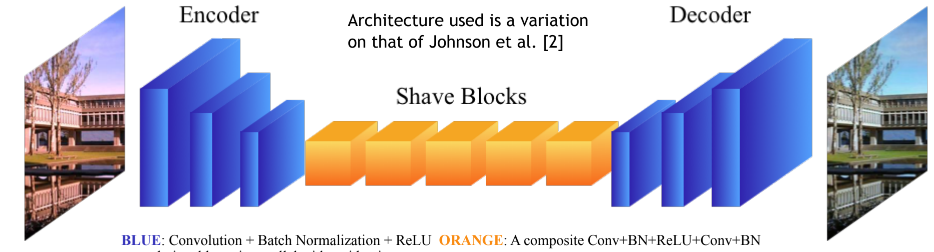

Architecture used is a variation on that of Johnson et al. [2]

BLUE: Convolution + Batch Normalization + ReLU **ORANGE**: A composite Conv+BN+ReLU+Conv+BN convolutional layer in parallel with an identity map.

Encode

type	kernel	stride	output
conv	9x9	1x1	32
conv	3x3	3x3	64
conv	3x3	3x3	128

Shave Block

type	kernel	stride	output
conv	3x3	1x1	128
conv	3x3	1x1	128
shave	-	-	128
sum	-	-	128

Decode

type	kernel	stride	output
deconv	3x3	2x2	64
deconv	3x3	2x2	32
conv	9x9	1x1	3

Kernel indicates the size of the convolution (conv) kernel. ***Stride*** controls the subsampling of the input data. ***Output*** refers to the number of convolution filters of the given size and stride used in the respective convolutional layer.

Example Results

Training and Test Sets

* Large datasets of images under spatially-varying illumination do not exist

* Synthesized applying spatially-varying scaling (von Kries) to R and G channels.

* Used COCO images. 50,000 for training, 10,000 for testing

* Linear variation across image
  For example –>

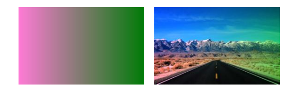

Median Angular Error Over All Pixels

Dataset	Median Input -> Recovered
NUS Canon	7.9 -> **6.1**
NUS Fujifilm	7.3 -> **4.8**
NUS Nikon	7.2 -> **4.4**
NUS Samsung	7.0 -> **3.8**
NUS Sony	7.2 -> **4.7**
MS COCO	7.0 -> **3.9**

Conclusion

* Significantly reduces spatially-varying color casts

* End-to-End processing

* Eliminates traditional illumination-estimation step

References

[1] S. Iizuka, E. Simo-Serra, and H. Ishikawa. Let there be color!: Joint end-to-end learning of global and local image priors for automatic image colorization with simultaneous classification. Proc. of SIGGRAPH 2016, 35(4):110:1–110:11, 2016.

[2] J. Johnson, A. Alahi, and Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. ECCV 2016

Contact

Ligeng Zhu: lykenz@sfu.ca
Brian Funt: funt@sfu.ca

For source code and paper scan ==>