What Else Can Fool Deep Learning?
Addressing Color Constancy Errors on Deep Neural Network Performance

Mahmoud Afifi[1] Michael S. Brown[1,2]

[1]York University [2]Samsung AI Center

ICCV 2019 Seoul, Korea

SCAN ME
Project page

SCAN ME
Source code

Introduction

Camera Imaging Pipeline

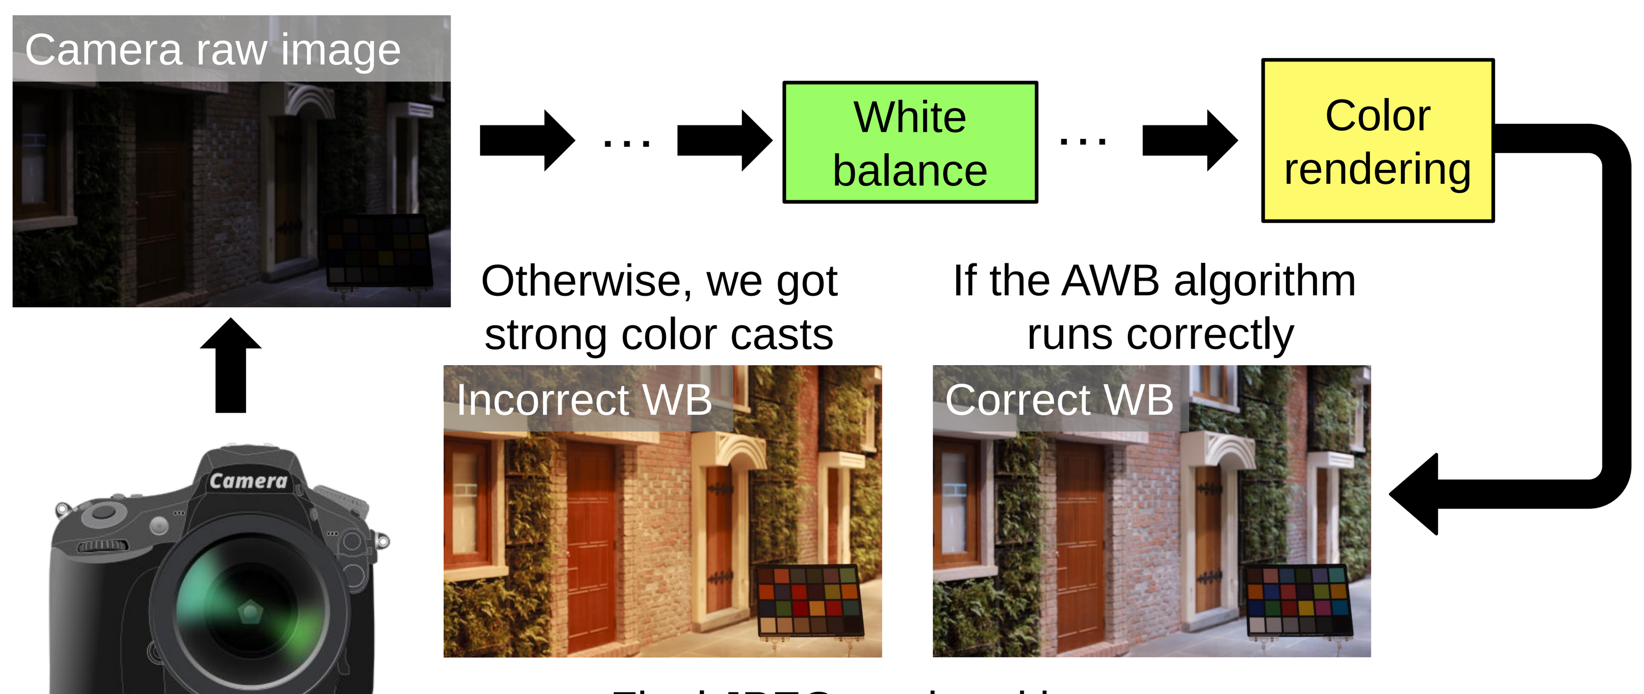

Final JPEG rendered images

Can such global color manipulations fool trained DNNs?

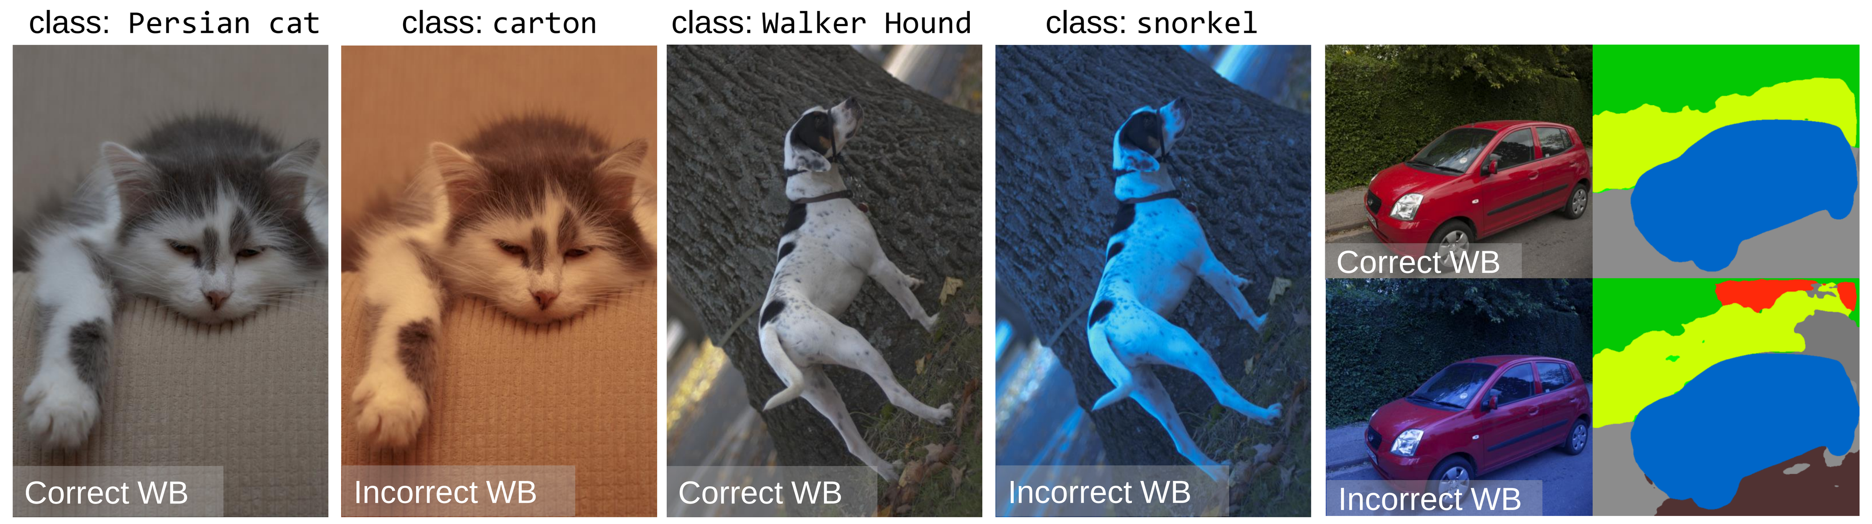

Contributions

- Examining how errors related to computational color constancy can *adversely* affect DNNs trained for ***image classification*** and ***image semantic segmentation.***
- Showing that traditional color augmentation methods are not well suited to deal with color errors caused by cameras.
- We propose a novel **color augmenter to emulate WB errors**.
- Our WB augmenter improves accuracy for image classification and semantic segmentation with **< 7.3 sec (CPU)** and **1.0 sec (GPU)** to generate **ten** full-resolution ~**12 Mpix**.

Our WB emulator

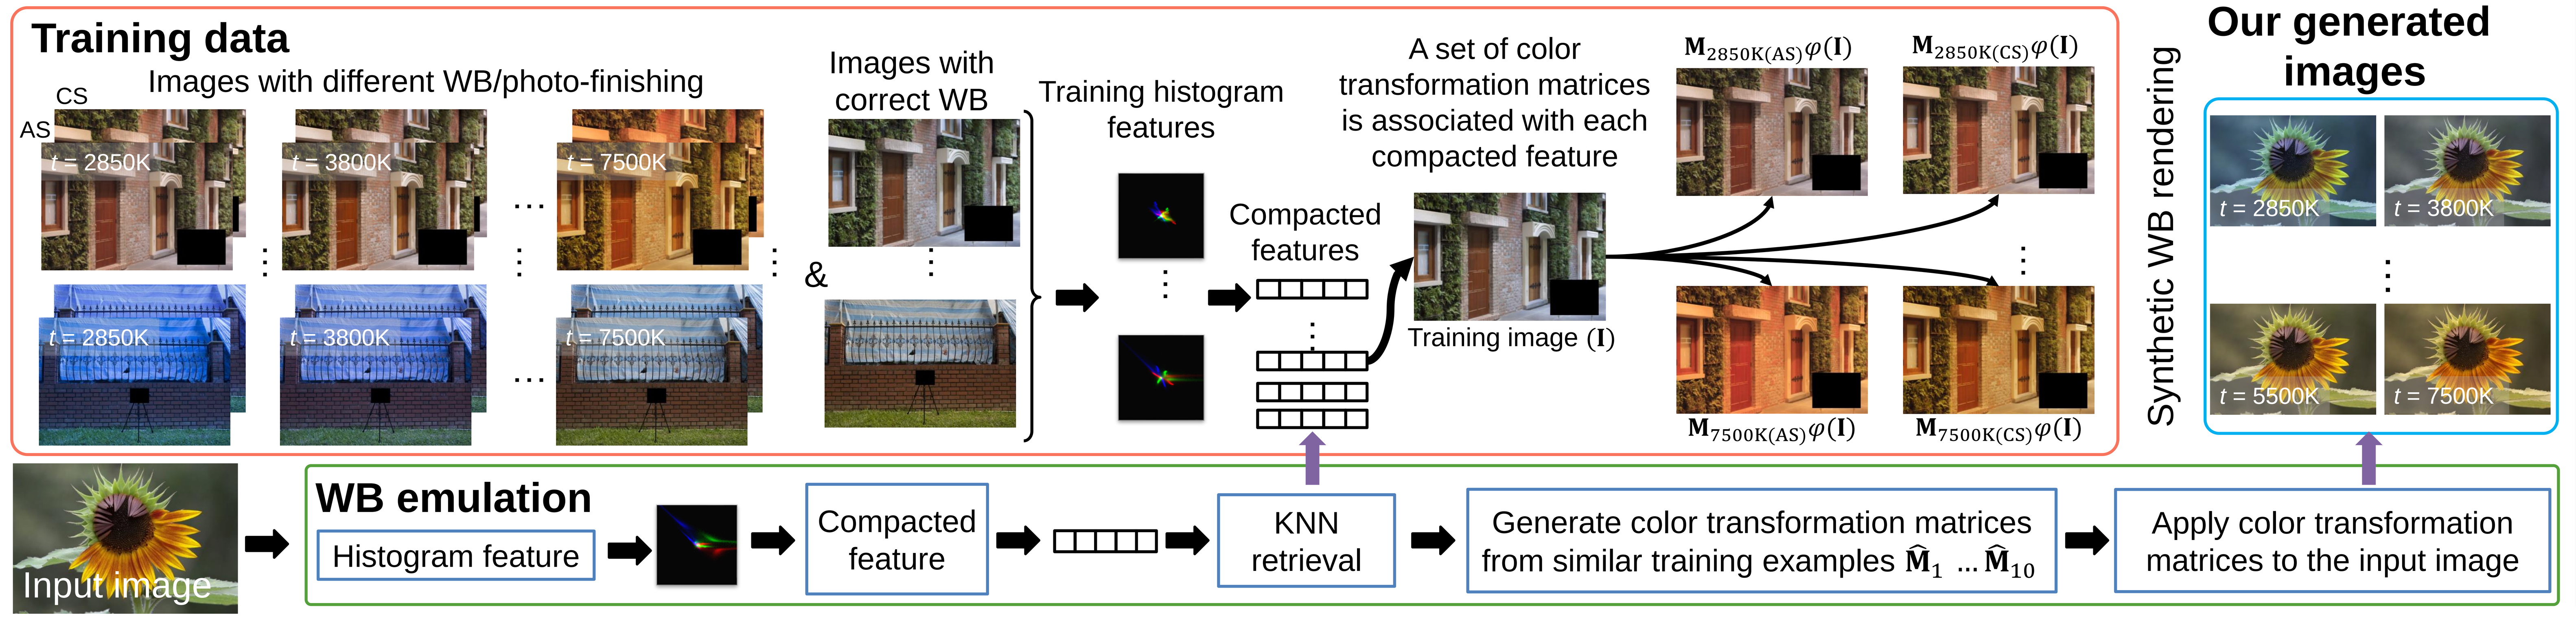

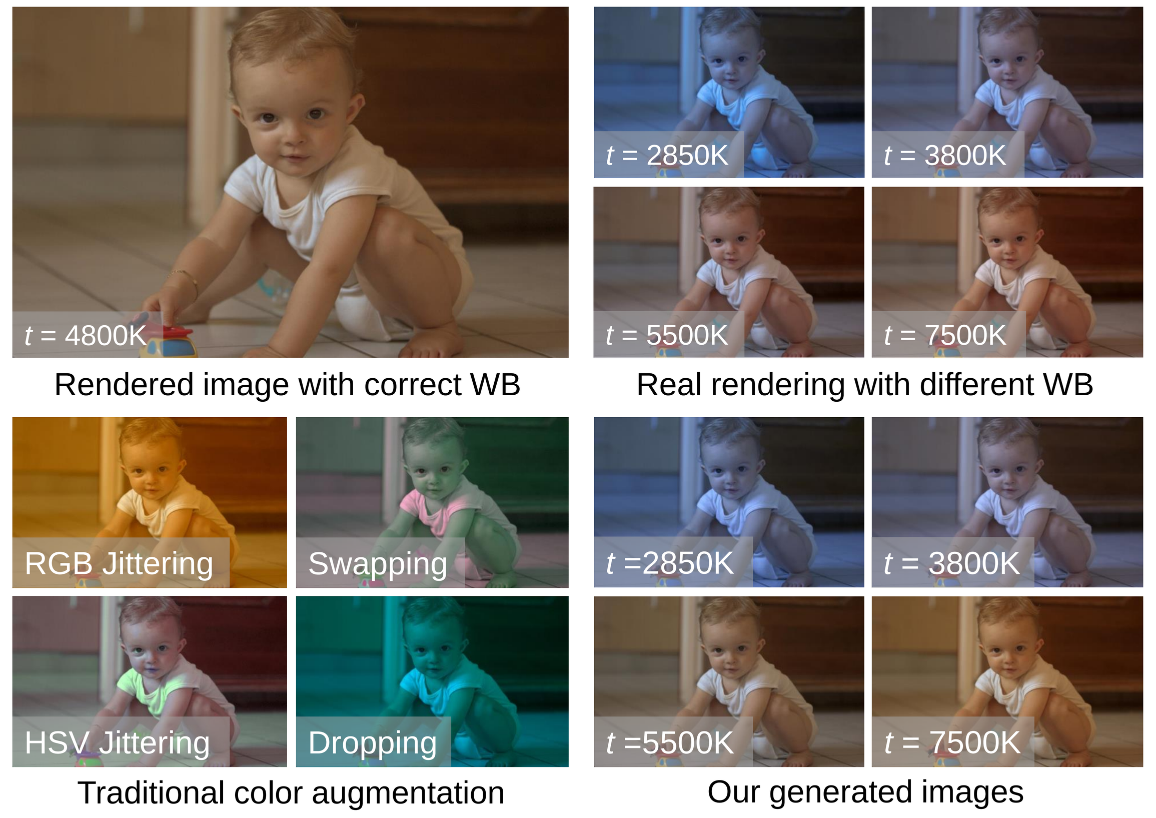

Effects of WB Errors on Pre-trained DNNs

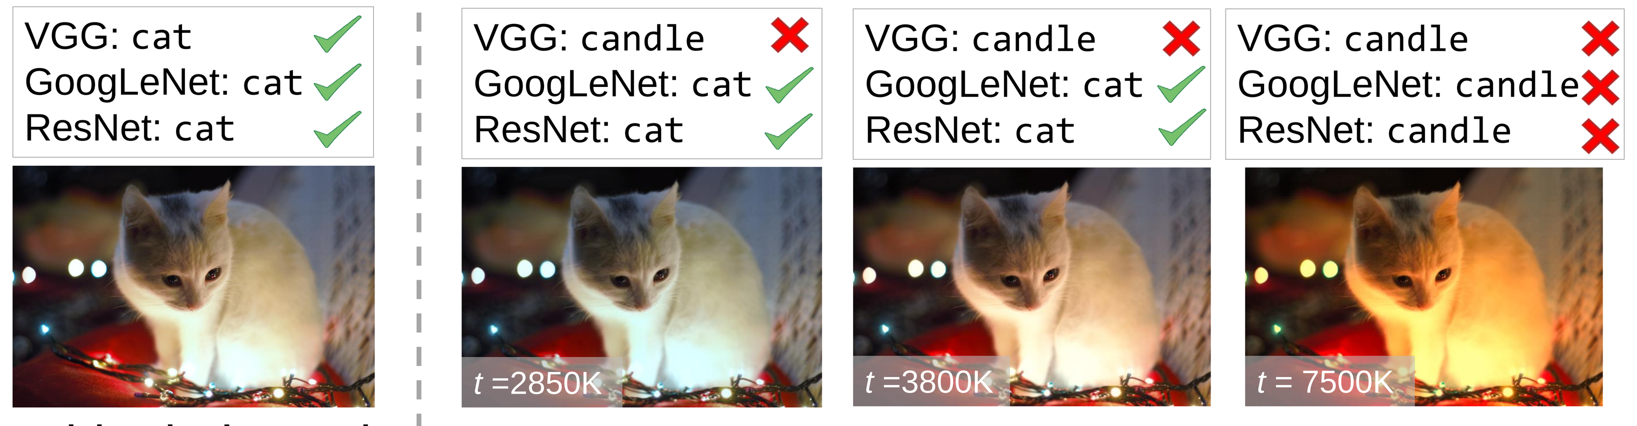

White-balanced image

Three generated images with different WB settings

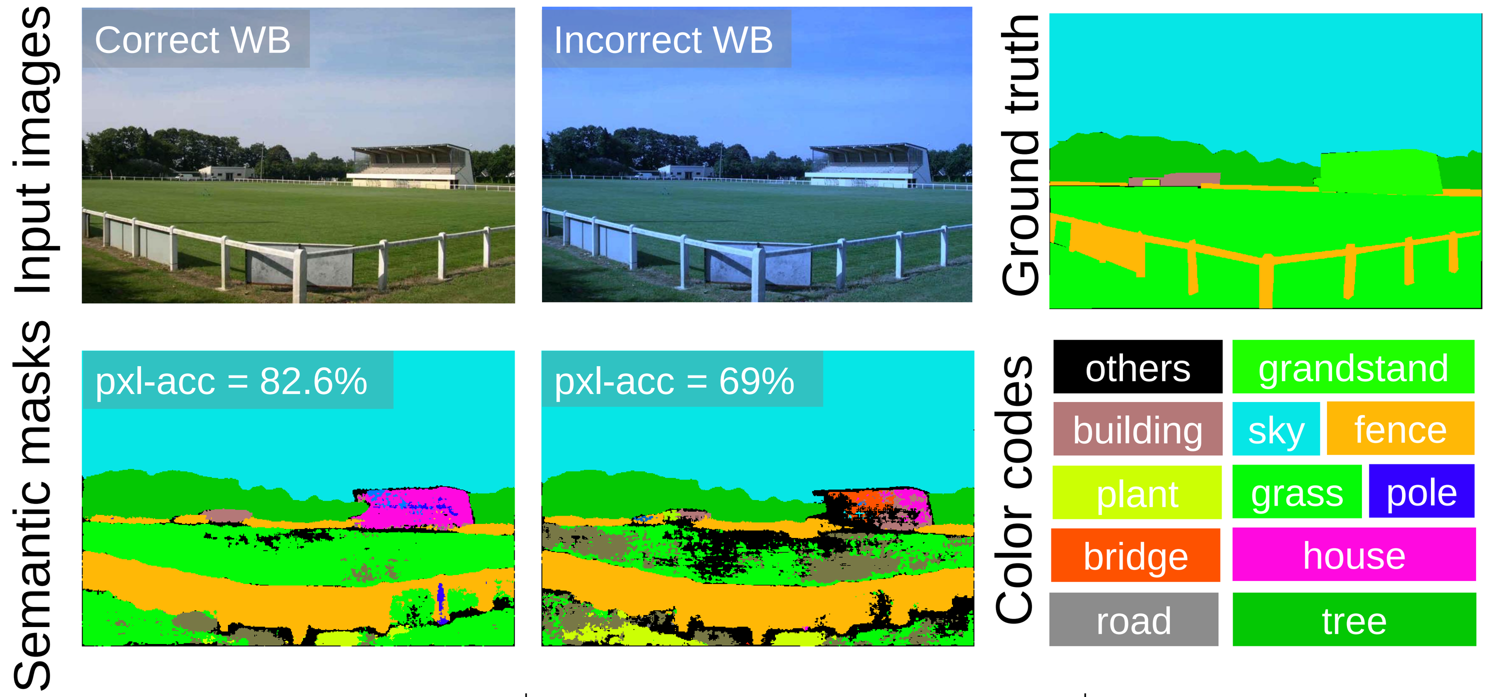

Image classification

ImageNet

AlexNet
VGG
GoogLeNet
ResNet

~-10%

Semantic segmentation

ADE20K

DilatedNet
RefineNet

~-2%

Evaluation

Proposed training strategies

- Train w/wo color augmentation (ours, other color augmenters)
  - We use the same number of iterations for fair comparisons
- Pre-processing WB correction

Datasets and models

- Image classification:

  CIFAR-10, CIFAR-100, ~15K images w different WB | SmallNet **(S)**, AlexNet **(A)**
- Semantic segmentation:

  ADE20K | SegNet

Testing set	No color aug.	HSV jittering aug.	RGB Jittering aug.	Our WB augmenter
Image classification (classification acc.)				
CIFAR-10	S: 0.80 | **A: 0.93**	S: 0.80 | A: 0.92 (↓ 0-1%)	S: 0.78 | A: 0.92 (↓ 1-2%)	**S: 0.81 | A: 0.93 (↑ 0-1%)**
CIFAR-10 (diff. WB)	S: 0.66 | A: 0.80 (↓ ~13%)	S: 0.75 | A: 0.86 (↑ 6-9%)	S: 0.77 | A: 0.87 (↑ 7-11%)	**S: 0.79 | A: 0.89 (↑ 9-13%)**
CIFAR-100	**A: 0.77**	A: 0.72 (↓ 5%)	A: 0.72 (↓ 5%)	A: 0.74 (↓ 4%)
CIFAR-100 (diff. WB)	A: 0.53 (↓ 24%)	A: 0.61 (↑ 8%)	A: 0.65 (↑ 12%)	**A: 0.67 (↑ 14%)**
Diff. WB + pre-WB	S: 0.49 | A: 0.77	S: 0.48 | A: 0.78 (↑ 1%)	S: 0.47 | A: 0.78 (↑ 1-2%)	**S: 0.48 | A: 0.79 (↓↑ 1.5%)**
Diff. WB	S: 0.40 | A: 0.73 (↓ 4-9%)	S: 0.46 | A: 0.77 (↑ 4-6%)	S: 0.48 | A: 0.78 (↑ 5-8%)	**S: 0.50 | A: 0.79 (↑ 5-10%)**
Semantic segmentation (pxl-acc)				
ADE20K	**0.60**	0.58 (↓ 2%)	0.54 (↓ 6%)	**0.60**
ADE20K (diff. WB)	0.56 (↓ 4%)	0.54 (↑ 2%)	0.53 (↓ 3%)	**0.58 (↑ 2%)**

References:

ImageNet: J Deng, et al., Imagenet: A large-scale hierarchical image database. In CVPR, 2009.
ADE20K: B Zhou, et al., Scene parsing through ADE20K dataset. In CVPR, 2017.
CIFAR: A Krizhevsky and G Hinton, Learning multiple layers of features from tiny images. Technical report, 2009.
SmallNet: L Perez and J Wang. The effectiveness of data augmentation in image classification using deep learning. arXiv preprint, 2017.
AlexNet: A Krizhevsky, et al., Imagenet classification with deep convolutional neural networks. In NIPS, 2012
ResNet: K He, et al., Deep residual learning for image recognition. In CVPR, 2016.
VGG: K Simonyan and A Zisserman, Very deep convolutional networks for large-scale image recognition. arXiv preprint, 2014.
GoogLeNet: C Szegedy, et al., Going deeper with convolutions. In CVPR, 2015.
DilatedNet: F Yu and V Koltun, Multi-scale context aggregation by dilated convolutions. In ICLR, 2015.
RefineNet: G Lin, et al., Refinenet: Multi-path refinement networks for highresolution semantic segmentation. In CVPR, 2017.
SegNet: V Badrinarayanan, et al., Segnet: A deep convolutional encoder-decoder architecture for image segmentation. TPAMI, 2017.
WB Correction: M Afifi, et al., When color constancy goes wrong: Correcting improperly white-balanced images. In CVPR, 2019.