Learning to Detect Multiple Photographic Defects

Ning Yu[1], Xiaohui Shen[2], Zhe Lin[2], Radomír Měch[2], and Connelly Barnes[1,2]
[1]University of Virginia, [2]Adobe Research

Code and data: https://github.com/ningyu1991/DefectDetection

Defect Definition and Severity Ranking

Goal

Simultaneously detect the existence and severity of seven photographic defects.

A Novel Dataset

12,583 natural images from Flickr

7 defects per image

5 annotation votes per defect

↓

449,855 annotations

Defect Severity Prediction

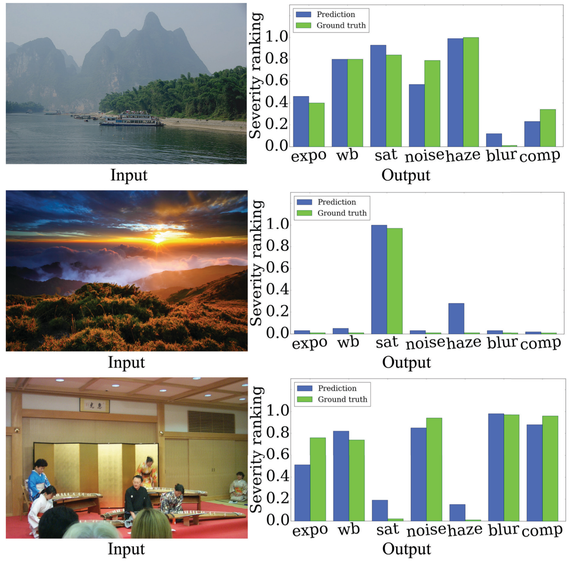

Ablation Study

	Mean Ranking Correlation
Single-Column	0.7646
Separate Networks	0.7788
Regression L^2 Loss	0.7879
Classification Loss	0.7740
Without Augmentation	0.7336
Multi-Column (Holistic)	0.7482
Multi-Column (Patch)	0.7789
Multi-Column (Combined)	0.8001

Comparison

	Noise	Haze	Undesired Blur
Previous	0.4199 [1]	0.6615 [2]	0.4864 [3]
Ours	0.8174	0.8490	0.6867

Network Architecture

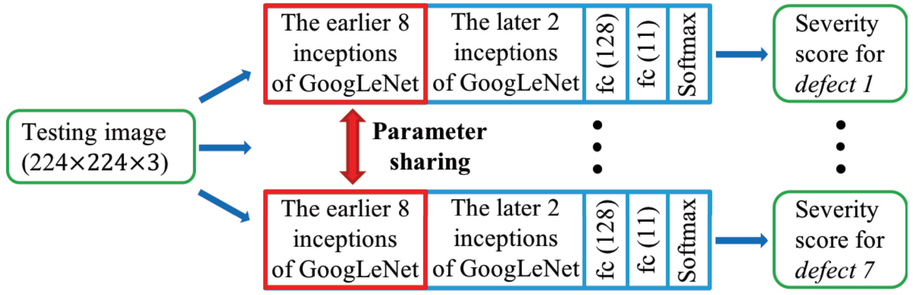

Defect Heatmap

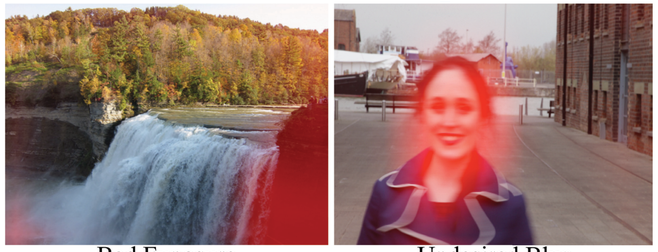

Reference
[1] X. Liu, M. Tanaka, and M. Okutomi. Single-image noise level estimation for blind denoising. In IEEE TIP 2013.
[2] K. He, J. Sun, and X. Tang. Single image haze removal using dark channel prior. In IEEE TPAMI 2011.
[3] A. Chakrabarti, T. Zickler, and W. T. Freeman. Analyzing spatially-varying blur. In CVPR 2010.