

Noise Flow: Noise Modeling with Conditional Normalizing Flows

—Supplemental Material—

Abdelrahman Abdelhamed^{1,2}
¹York University

Marcus A. Brubaker²
²Borealis AI

Michael S. Brown^{1,3}
³Samsung AI Center, Toronto

{kamel, mbrown}@eecs.yorku.ca, marcus.brubaker@borealisai.com

In this supplemental document, we elaborate more on the application of Noise Flow for training convolutional neural network (CNN) image denoisers and present more generated noise samples from Noise Flow for visual inspection.

Denoising Training Loss

Figure 1 shows the training loss over 2000 epochs of the DnCNN [1] image denoiser using the three noise synthesis models: (1) DnCNN-Gauss: the Gaussian model; (2) DnCNN-CamNLF: the heteroscedastic Gaussian model represented by the camera-calibrated noise level functions (NLFs); and (3) DnCNN-NoiseFlow: our Noise Flow model. Despite that training loss of the DnCNN-CamNLF is the lowest, the training behaviour indicated that DnCNN-NoiseFlow is the most stable. The DnCNN-Gauss model is also stable, however, it still yields lower testing performance, as shown in Figure 2.

Denoising Testing Results

Figure 2 shows the testing peak signal-to-noise ratio (PSNR) over 2000 epochs of the three models from Figure 1. DnCNN-CamNLF outperforms the DnCNN-Gauss models by a small margin, despite the lower training loss of the former. The DnCNN-NoiseFlow model yields the best performance despite having slightly higher training loss than DnCNN-CamNLF.

Figure 3 shows more denoising results from the three models for visual inspection. The samples confirm the better performance of the DnCNN-NoiseFlow model and the importance of having more accurate noise models for generating realistic synthetic noise.

More Generated Noise Samples

Figure 4 shows more generated noise samples from the Noise Flow model compared to the Gaussian and the camera NLF models. Noise Flow samples are much closer to the real samples in terms of the marginal KL divergence. We show the corresponding ISO level and lighting condition on the left.

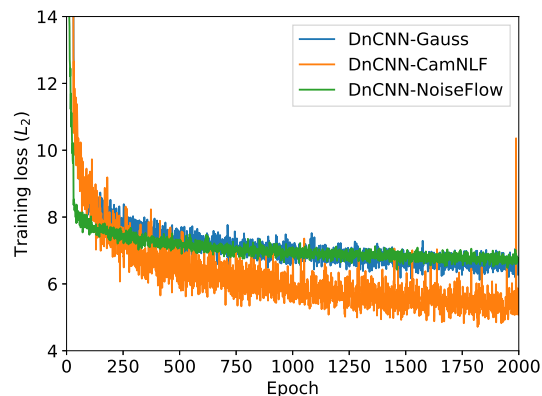


Figure 1: Results of training the DnCNN [1] image denoiser with synthetic noise generated by: the Gaussian model; the camera NLFs; and our Noise Flow model.

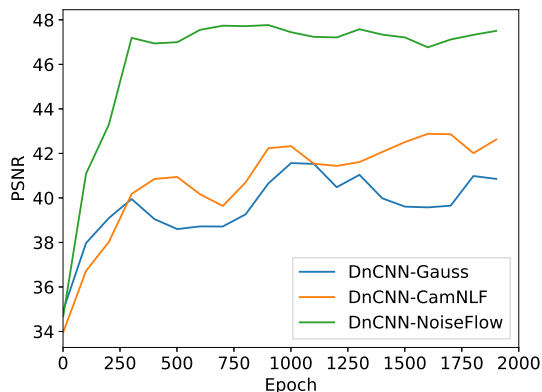


Figure 2: Testing PSNR results corresponding to the three model from Figure 1. The DnCNN model trained on noise generated with Noise Flow yields the best PSNRs.

References

- [1] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang. Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *TIP*, 2017. 1



Figure 3: Sample denoising results from the DnCNN denoiser trained on three different noise synthesis methods:(b) Gaussian; (c) camera NLF; and (d) Noise Flow. (a) Real noisy image. (e) Ground truth.

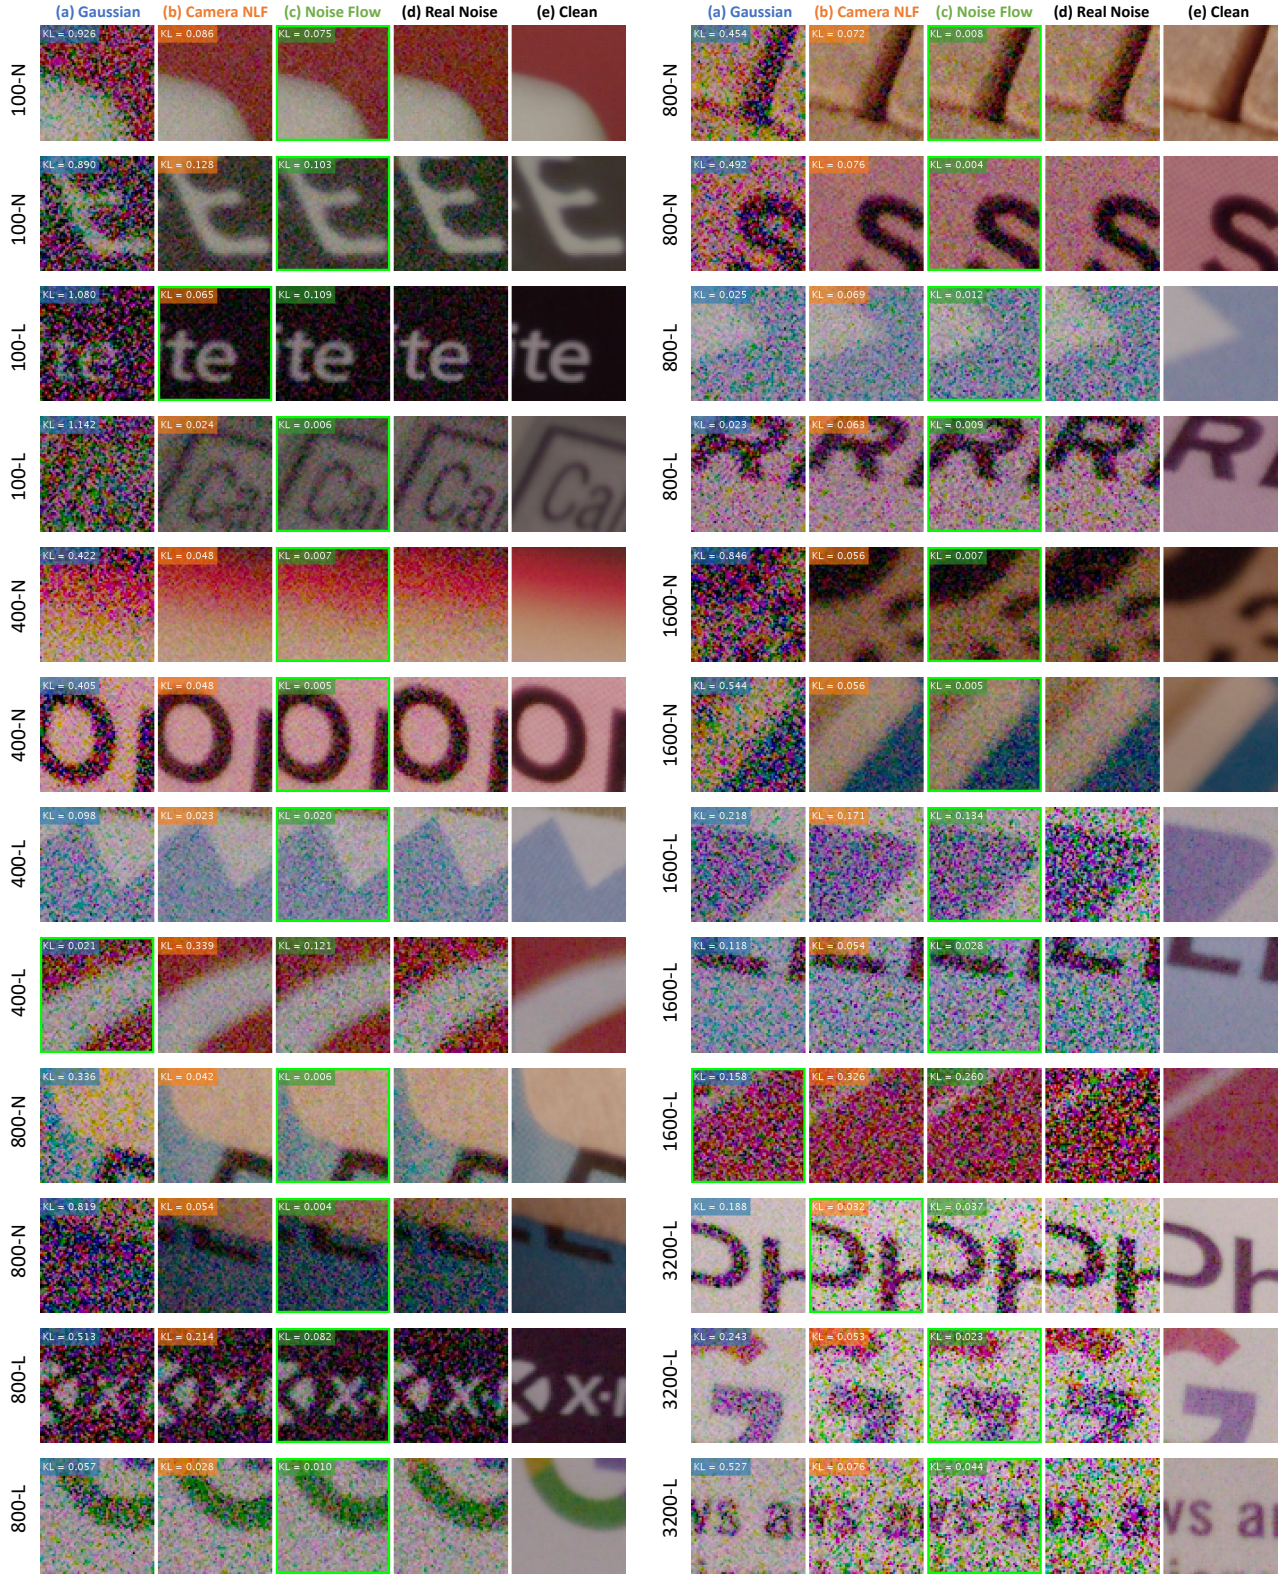


Figure 4: More generated noise samples from (c) Noise Flow compared to samples from (a) the Gaussian and (b) camera NLF models. (d) Real noise. (e) Clean image. Noise Flow samples are much closer to the real samples in terms of the marginal KL divergence. We show the corresponding ISO level and lighting condition on the left.