Image Fixed Pattern Noise Correction

Team: Paul Lim (paulglim@stanford.edu),
Steve Mims (mimsy@stanford.edu)

Problem Statement

- Photon shot noise can hide fixed pattern noise, if the ratio is 10:1 or better. Our customers require 20:1. In low light, with little shot noise, uncorrected FPN can be higher than the signal.
- Image correction can be done by calibrating individual die, but some customers do not have a budget for this time-consuming step.
- We need a universal FPN corrector, good for all die, and are training an 11-layer convolutional network to do that.

Network

- Following the approach of He, Cao, Dong et al ("Single-image-based nonuniformity correction..") Applied Optics, Vol 57,No 18, 20-June-2018)
- Convolution with a 5x5 kernel (to span four sequential row types), with 32 filters.

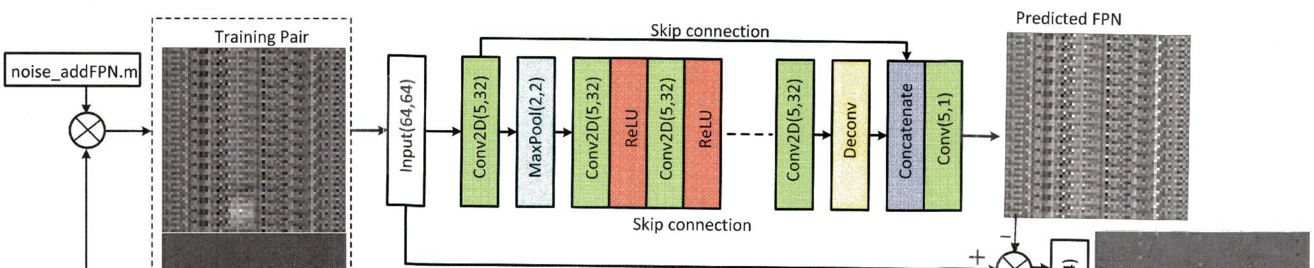

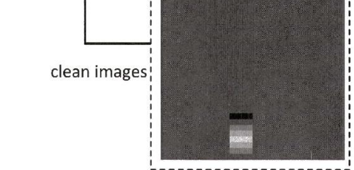

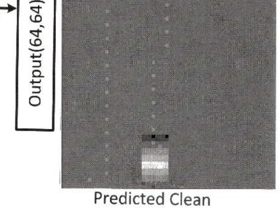

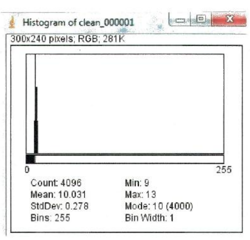

Training

Loss Function:

$$J = \frac{1}{MN} \sum_{i \in M, j \in N} |P_{i,j} - T_{i,j}|$$

where M,N are the number of rows and columns of the input images respectively, and P is the predicted image output and T is the original clean image output

Loss function	final learning rate	epochs	final training loss	final validation loss
mean_abs_error mean(abs(yhat-y))	5.00E-07	3660	**8.30E-04**	**8.29E-04**
mean_abs_percent_error mean(abs(y-yhat) /max(y,epsilon))*100	2.00E-06	130	5.36E+03	5.51E+03

He, Cao, et al suggested L1 norm (mean_abs_error) produced more human-pleasing images than L2 norm. We also tried mean_abs_percentage_error, but discarded. We think having too many zero-value pixels caused an over-emphasis on these. Started w/learning_rate=1e-4. Final training/validation loss was 8.3e-4, equivalent to 0.21 DN. Adam optimizer with momentum was used with minibatches of 64.

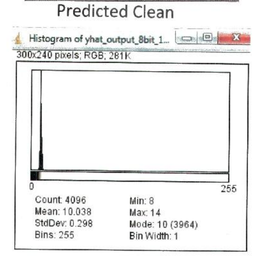

Training Data

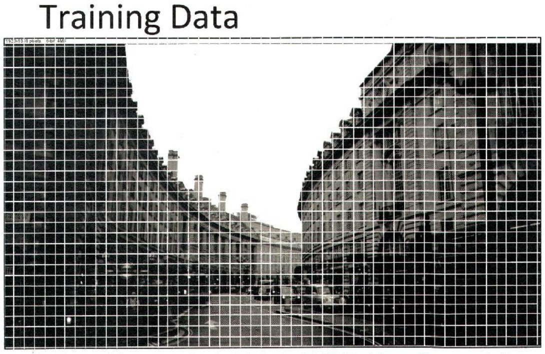

100 8-bit HD grayscale images were downloaded, sliced up into 64x64 sub-images, then augmented by rotation and flipping, to create 100,000 "clean" sub-images. Then, Matlab code was run to add random fixed-pattern-noise (horizontal and vertical) to these clean images to create 100,000 noisy images. The noise was based on approximately 40,000 parameters (e.g. gain/offset), collected on a 4K image sensor from BAE Systems.

Discussion and Future work

- Using percentage error for loss might have emphasized smaller pixel values as we wished, but would require more work with dataset to remove zero values.
- 120x speed improvement with AWS makes huge difference in minimum error that can be achieved.
- Expand to larger image sizes (e.g. from 4K sensors)
- Debayering for colored sensors.

References:

He, et al; "Single-image-based nonuniformity correction of uncooled long-wave infrared detectors: a deep learning approach"; Applied Optics, Vol 57 No. 18; June 20 2018, D155