A Deep Learning Approach to JPEG Colour Restoration Without Ground Truth

Michael Barrack
Master's In Applied Computing (MScAC), University of Toronto

Motivation

- JPEG compression uses **Chroma Subsampling** to reduce the complexity of colour content, which can significantly degrade colour accuracy, especially at lower quality factors.

Original Image   JPEG, Quality = 10

- In recent years there have been investigations into image denoising techniques using Deep Convolutional Neural Networks (CNNs) that do not require the availability clean ground-truth data [1][3]
- The goal of this project is to evaluate the effectiveness of these approaches at recovering the colour content lost during Low-Quality JPEG compression, and to experiment with novel adaptations that may improve their performance in this area

Methodology

The training structure follows from that used in the related works. A CNN is initialized randomly and optimized to reconstruct a single JPEG from either uniform noise, or itself. [3] [4]

Comparative testing of the model structures included: compression quality, model complexity, and the use of skip connections.

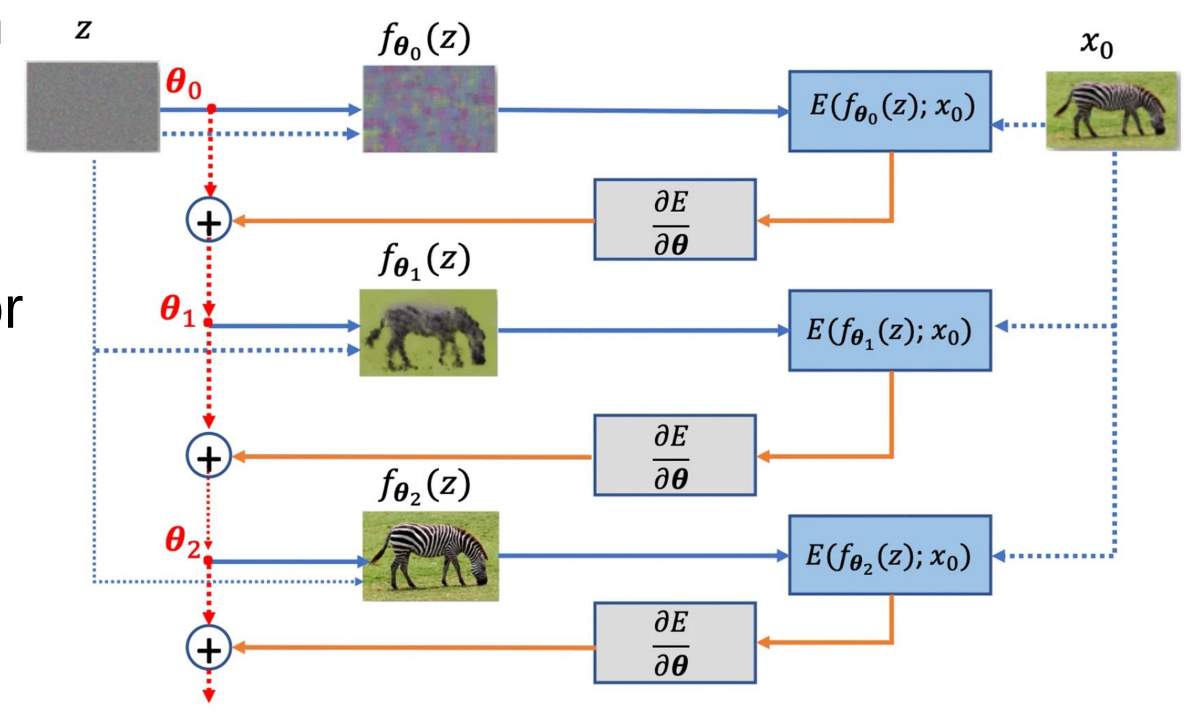

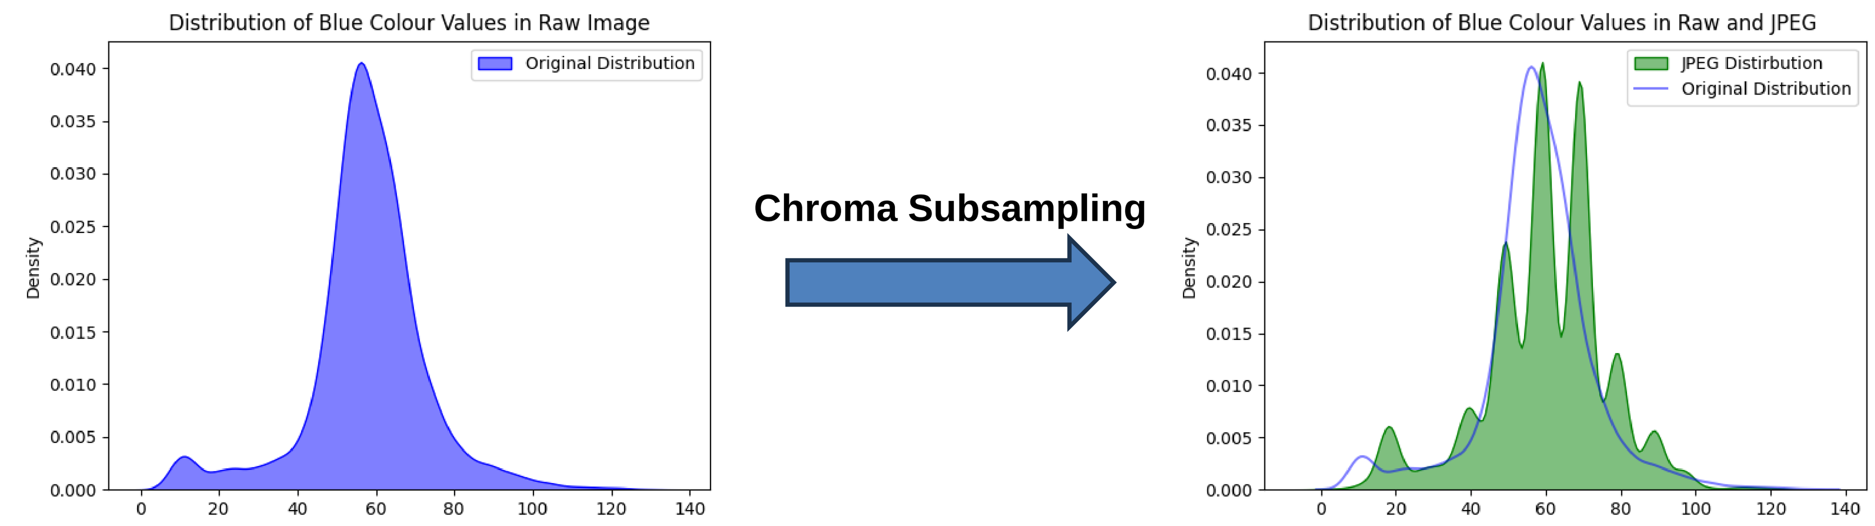

Based on the observed change in distribution of colour values, it was decided to equip these network structures from *Deep Image Prior* and *Unsupervised Learning with SURE,* with a new loss function:

$$Loss(x, \hat{x}) = MSE(x, \hat{x}) + \frac{\lambda}{3} \sum_{C \epsilon \{R,G,B\}} D_{KL}(P_{x_C} || Q_{\widehat{x_C}})$$

Here, D_{KL} is the Kullback-Liebler Divergence. For each colour band, P_{x_C} is a is a Gaussian distribution defined to have the same mean and variance as the JPEG's distribution, and $Q_{\widehat{x_C}}$ is the distribution in the reconstructed image. The average KL-Divergence value across the three colour bands is weighted against mean-squared error loss by λ. The idea here is to encourage the network to smooth the distribution of colour values.

Related Work

- **Learning Image Restoration without Clean Data [1]:** Presents a scheme where a noisy dataset is further distorted, and a model is trained to map these corrupted images to their counterparts.
- **Deep Image Prior [3]:** This paper focuses on the ability of the structure of a CNN (with and without skip connections) to serve as a prior for image restoration tasks, without requiring any external training data.
- **Unsupervised Learning with Stein's Unbiased Risk Estimator (SURE) [4]:** Building on the work in *Deep Image Prior*, this paper explores the use of SURE as a loss function to prevent overfitting by providing an "*unbiased estimate of the MSE"*, improving model performance without requiring clean ground-truth data.

Experimental Results

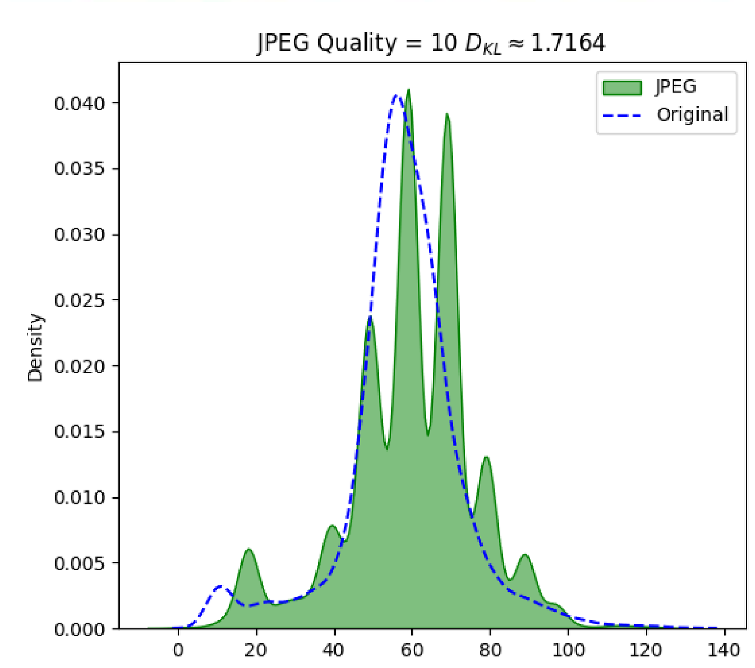

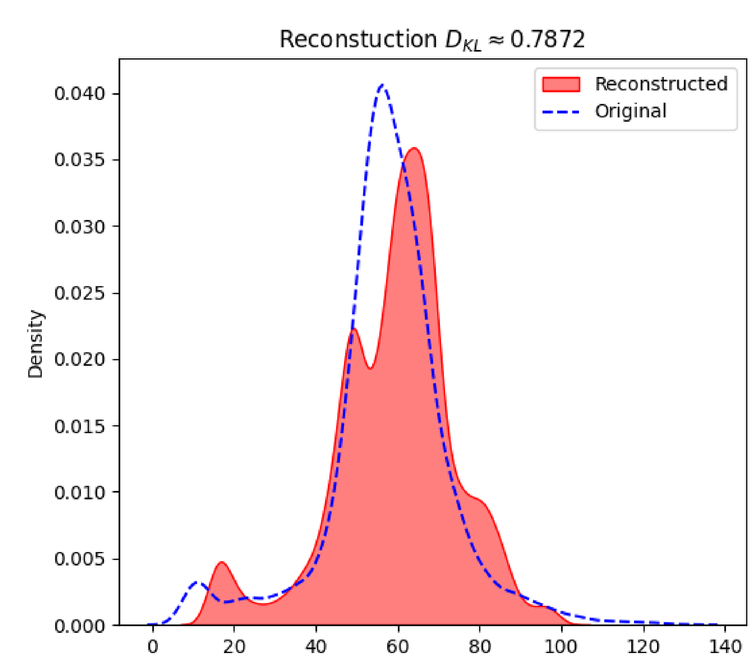

Model Average PSNR (dB)

	JPEG Quality	*Deep Image Prior (MSE)*	*SURE*	*Weighted KL-Divergence λ = 10*	*Weighted KL-Divergence λ = 100*
WRT JPEG	***10***	31.1	30.42	31.02	31.07
	15	29.7	30.45	30.93	30.79
	25	29.6	30.40	30.67	30.74
WRT Ground Truth	***10***	26.8	26.45	26.76	26.75
	15	27.5	27.45	27.84	27.78
	25	28.3	28.33	28.78	28.85

References

[1] Lehtinen et al. Noise2noise: Learning image restoration without clean data. CoRR, abs/1803.04189, 2018.
[2] Mao et al. Image restoration using convolutional auto-encoders with symmetric skip connections. CoRR, abs/1606.08921, 2016.
[3] Ulyanov et al. Deep image prior. CoRR, abs/1711.10925, 2017
[4] Metzler et al. Unsupervised learning with stein's unbiased risk estimator. CoRR, 2020.