



The Impact of JPEG Compression on Prior Image Noise

Marina Gardella, Tina Nikoukhah, Yanhao Li, Quentin Bammey

► To cite this version:

Marina Gardella, Tina Nikoukhah, Yanhao Li, Quentin Bammey. The Impact of JPEG Compression on Prior Image Noise. 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP 2022), May 2022, Singapour, Singapore. pp.2689-2693, 10.1109/ICASSP43922.2022.9746060 . hal-03907095

HAL Id: hal-03907095

<https://hal.science/hal-03907095v1>

Submitted on 19 Dec 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THE IMPACT OF JPEG COMPRESSION ON PRIOR IMAGE NOISE

Marina Gardella Tina Nikoukhah Yanhao Li Quentin Bammey

Université Paris-Saclay, Université Paris Cité, ENS Paris-Saclay, CNRS, Centre Borelli

ABSTRACT

JPEG compression is widely used to store digital images and extensive studies analysed its impact on the image quality; in particular the quantization noise and artefacts created by JPEG. Nevertheless, there is little work on the impact of JPEG compression on the noise already present in the image. In this paper, we propose a model predicting how the noise power is affected by JPEG compression. This allows for a better understanding the noise traces on the image, which is crucial for image forensic analysis and image restoration. An interactive demo for this article is available at

<https://ipolcore.ipol.im/demo/clientApp/demo.html?id=77777000136>

Index Terms— JPEG compression, camera noise, noise estimation, DCT analysis, quantization

1. INTRODUCTION

From the moment a raw image is acquired until the final JPEG picture is obtained, a complex processing chain is applied. Each of these operations alters the noise model. Indeed, the initial Poisson-Gaussian noise [1] undergoes several operations resulting in a complex noise model. The final stage in most digital images consists in JPEG compression. Indeed, to be stored or transferred in a reasonable amount of time, images must undergo a compression step. Many noise estimation algorithms [2]–[5] suppose that noise can be estimated using only the high frequency coefficients of small patches in an image. However, this is not true for JPEG-compressed images, since the quantization step during the compression process attenuates these high frequencies. Indeed, for JPEG images, noise variance decreases as the frequency increases. As a result, these noise estimation methods yield to an inaccurate estimation of the noise. Noise estimation is a mandatory step of countless image processing tasks such as denoising [6], [7], forgery detection [8], [9], anomaly detection [10], PRNU extraction [11] and steganography [12], just to mention a few. The performance achieved by the methods developed to tackle each of these tasks depend on how accurately they are able to estimate noise.

Work funded by French *Ministère des Armées – Direction Générale de l’Armement* and *Région Île-de-France*.

Centre Borelli is also a member of Université Paris Cité, SSA and INSERM.




Fig. 1: Simplified processing pipeline of an image, from its acquisition by the camera sensor to its storage as a JPEG-compressed image. The right column plots the noise of the image as a function of intensity in all three channels.

The aim of this article is to provide a characterisation of the resulting noise after JPEG compression. Such a model could help accurately estimate noise, boosting the performance of a huge variety of image processing tasks that require noise estimation as one of their fundamental steps.

2. RELATED WORK

After JPEG compression, there are two kinds of noise in the image: the original noise, that was there before compression and was compressed with the rest of the image, and the noise coming from the JPEG compression itself. There is consequent literature on estimating the noise that directly comes from compression itself [13]–[15]. Due to the lossy nature of JPEG, it affects tasks in various domains. In image forensics, subtle traces such as demosaicing artefacts are much harder to analyse when the image is compressed [16]. Mandelli, Bonet-




Fig. 2: Gaussian noise is added to the `dice` image before compression. An estimation of the remaining noise is obtained by computing the difference between the noisy and noiseless images, both after compression. This enables us to estimate specifically the effect of the compression on existing noise, while ignoring most of the noise coming directly from the compression. As per our model, a low noise level is diminished even more by the compression, whereas a higher noise level is instead augmented. Compression is done with the Pillow library at a quality factor of 85.

tini, Bestagini, *et al.* [17] analyses the effects of JPEG compression on different camera traces and shows that care must be taken from the training data to ensure some robustness to compression. Image classification, especially when done at full resolution, is also affected by compression [18].

In comparison, analysis of how the compression affects already-present noise has received little attention. We propose to study in details how JPEG compression affects noise that was present prior to the compression. Indeed, noise is present in an image from the sensing of the real scene. Each step in the pipeline, including JPEG compression, alters the already-existing noise in some way, while often adding its own noise on top of it. See [19] for a general overview of the impact of each processing step on the noise. Jiang, Zeng, Kang, *et al.* [20] notice that counter-forensics techniques, aimed at hiding JPEG compression, usually introduce inconsistent noise in the image. Noise-level analysis is applied to distinguish authentic regions from region with hidden traces of a previous compression. Corchs, Gasparini, and Schettini [21] study the influence of images distortions, in particular Gaussian noise and JPEG compression, on the overall image quality. While this study combines Gaussian noise that is subjected to further JPEG compression, they do not directly study the influence JPEG compression had on the noise, but rather the combined effect both had on the end image’s quality.

3. NOISE THROUGHOUT THE CHAIN

Along the image formation pipeline of a camera, the raw data from the sensor undergoes a series of operations: demosaicing, white balance, gamma correction, compression, to mention a few [22]. We will now discuss the way in which noise is affected at each step of the camera processing chain.

In order to acquire a raw image, the first step is to count the number of photons impacting over the CFA during the ex-

posure time. This is done by transforming the incoming light photons into electronic charge, which is stored in a potential well and then turned into analog voltage outputs. The final step consists in the conversion of the analog voltage measures into digital quantized values. The value at each pixel at this stage can be modelled as a Poisson variable whose expectation is the real pixel value. Furthermore, all channels have the same noise curve. Since noise is Poisson distributed, noise variance follows a simple linear relation as shown in Figure 1.

The raw image obtained from the CFA has one colour component per pixel, red, green, or blue. The demosaicing process consists in the reconstruction of a full colour image which is done by interpolating the two missing colour values per pixel. Figure 1 shows that, after demosaicing, each channel has a different noise curve. This is due to the fact that channels are processed differently by the demosaicing algorithm. After this, noise is spatially correlated.

The third step consists in performing colour correction so that the observed colours are true to the real scene. This is done by scaling the colour channels so that true gray objects are rendered as such in the final image. The pixels’ values are then tuned to accurately represent human vision, which is not linear with the signal intensity [23]. This is usually done by applying a power law function or by using pre-computed curves to fit the visual system. At this point, the noise is no longer monotonically increasing.

The final step for most digital images is JPEG compression. The colour space is switched from RGB to $YC_B C_R$. The chroma components C_B and C_R are usually downsampled. Each channel is then split into 8×8 blocks and the 2D DCT II of each block is computed. The resulting blocks are then quantized according to a quantization table, which is associated with the JPEG global quality and provides a different factor q for each component of the DCT blocks. It is during this quantization step that the greatest loss of informa-




Fig. 3: Ground truth residual computation from noiseless uncompressed images.

tion occurs, but it is also this step that allows the most space in memory to be saved. The coefficients that correspond to the high frequencies, whose variations the human visual system struggles to distinguish, are the most quantized, sometimes even entirely cancelled. After quantization, a lossless compression (Huffman encoding) is applied in order to store the resulting DCT values. This process is explained for the luminance channel in the next section.

4. JPEG COMPRESSION ON GAUSSIAN NOISE

Although the noise present in an image before JPEG compression is not Gaussian, the Poisson distribution of noise derived from the photon count can be approximated by a Gaussian distribution. However, this Gaussian distribution would still be intensity dependent. In order to transform this signal dependent noise into homoscedastic, a variance stabilizing transformation can be applied [24]. This procedure allows us to work with the popular white Gaussian noise assumption used in many noise estimation and denoising algorithms.

Let n be the spatially-independent zero-mean Gaussian noise existing in an image before JPEG compression. We note by $\tilde{u} \triangleq u + n$ the noisy observation in pixels, where u is the ideal noiseless image. Here we focus on the luminance colour space which is not downsampled and thus preserves more image details. Excluding the effect of lossless compression ($QF = 100$), the impact of lossy JPEG compression can be modelled as follows:

Consider an 8×8 block, the first operation consists in performing the DCT transformation of the block. Then, the DCT coefficients are converted into integers with division by a quantization table, followed by rounding. During the decoding stage, the multiplication by the quantization table is performed, followed by the inverse DCT (IDCT). The whole JPEG encoding-decoding process can be expressed as:

$\tilde{u}' = \text{IDCT} \circ \text{Quant} \circ \text{DCT}(\tilde{u}) = \mathbf{A}^\top [\mathbf{A} \tilde{u} \odot \frac{1}{q}] \odot q$, where $\mathbf{A} \in O(64, \mathbb{R})$ is the orthogonal matrix of DCT, $[\cdot]$ represents the rounding operation, $q \in \mathbb{R}_+^{64}$ is the quantization table and $\odot$ denotes the element-wise multiplication.

Since n is Gaussian, the noisy observation $\tilde{u}$ is in fact a Gaussian vector. Furthermore, if we assume that for all the pixels in a patch, the noise level is constant and equal to σ then $\tilde{u} \sim \mathcal{N}(u, \Sigma_{\tilde{u}})$, with $\Sigma_{\tilde{u}} = \sigma^2 I$.

The DCT is an orthogonal transformation and, therefore, $x \triangleq \text{DCT}(\tilde{u}) \sim \mathcal{N}(\mathbf{A}u, \Sigma_x)$ where $\Sigma_x = \Sigma_u = \sigma^2 I$.

The input of $[\cdot]$ in the pipeline, noted as $z \triangleq x \odot \frac{1}{q}$ is still a Gaussian vector: $z \sim \mathcal{N}(\mu_z, \Sigma_z)$ with $\mu_z = \mathbf{A}u \odot 1/q$ and $\Sigma_z = \text{diag}(\sigma^2/q_1^2, \dots, \sigma^2/q_{64}^2)$.

However, the impact of the rounding operation $[\cdot]$, described in detail in [25], is not linear. The rounded z is $z' \triangleq [z]$ has expected value for each entry

$$\mathbb{E}[z'_i] = \mu_{z,i} + \frac{1}{\pi} \sum_{l=1}^{\infty} e^{-\frac{2(\sigma\pi l)^2}{q_i^2}} \sin(2\pi l \mu_{z,i}) \frac{(-1)^l}{l} \quad (1)$$

Compared to textured patches, evaluating noise on flat patches is easier since the signal is uniform and the variation only contains noise. Thus we focus on the study of $\tilde{u} \leftarrow \tilde{u} - \text{mean}(\tilde{u})$ for the pre-compression image and $\tilde{u}' \leftarrow \tilde{u}' - \text{mean}(\tilde{u}')$ for the post-compression image. Then we get $\tilde{u} \sim \mathcal{N}(0, \sigma^2 I)$ and $x \sim \mathcal{N}(0, \sigma^2 I)$. The input of $\text{Quant}(\cdot)$ is $z \sim \mathcal{N}(0, \text{diag}(\sigma^2/q_1^2, \dots, \sigma^2/q_{64}^2))$ which has zero mean. The rounded zero-mean Gaussian variable has a zero expectation $\mathbb{E}[z'_i] = \mathbb{E}[z_i] = 0$ according to Eq. 1, and its variance is

$$\text{Var}(z'_i) = \frac{\sigma^2}{q_i^2} + \frac{1}{12} + \sum_{l=1}^{\infty} (-1)^l \left(\frac{1}{(\pi l)^2} + \frac{4\sigma^2}{q_i^2} \right) e^{-\frac{2(\sigma\pi l)^2}{q_i^2}} \quad (2)$$

Finally, the variation after compression $\tilde{u}'_i = \mathbf{A}_{:,i}^\top z' \times q_i$ satisfies $\mathbb{E}(\tilde{u}'_i) = 0$. If we define $\sigma_j'^2 \triangleq \text{Var}(\tilde{u}'_j)$, we get

$$\sigma_j'^2 = \sum_{i=1}^{64} (\mathbf{A}_{i,j})^2 \text{Var}(z'_i) q_i^2 \quad (3)$$

In particular, if the quantization table satisfies $q_i = q \forall i$ then,

$$\sigma_j'^2 = \sum_{i=1}^{64} (\mathbf{A}_{i,j})^2 \text{Var}(z') q^2 = q^2 \text{Var}(z') \triangleq \sigma'^2 \quad (4)$$

which is independent of the pixel index j . We get the relation between the standard deviation σ of the original noise and σ' that of the quantized noise:

$$\frac{\sigma'}{q} = \sqrt{\frac{\sigma^2}{q^2} + \frac{1}{12} + \sum_{l=1}^{\infty} (-1)^l \left(\frac{1}{(\pi l)^2} + \frac{4\sigma^2}{q^2} \right) e^{-\frac{2(\sigma\pi l)^2}{q^2}}} \quad (5)$$

5. EXPERIMENTS

To validate the model derived in Section 4 we conducted several experiments on both synthetic noise images and real noiseless images to which we added noise. The Noise-Free Test Images dataset [26] contains 16 high-quality images that were carefully downsampled to remove traces of previous noise, allowing us to control the amount of noise to add




Fig. 4: Noise residual standard deviation for each DCT coefficient, for quality factors $QF = 10, 50, 90$ and for a fixed pre-compression noise $\sigma = 10$. The top row corresponds to synthetic Gaussian noise (constructed by taking the noiseless uncompressed image as a completely flat image) and the bottom row to an image from the NFTI dataset with added noise [26].

to the image. To compute the ground truth noise residuals we first select an uncompressed noiseless image (flat or from the dataset). Given a quality factor QF and a noise level σ , the noiseless uncompressed image is processed in two different ways. The first one consists in adding white noise of variance σ^2 to the uncompressed noiseless image and then compressing it with quality factor QF . The second consists only in performing the compression step, without adding noise. The ground truth noise residual is computed as the difference between the noisy compressed image and the noiseless compressed image. This computation is summarised in Figure 3.

Figure 4 shows the standard deviation of the noise residual for each 8×8 DCT coefficient and for different compression qualities, and for a fixed pre-compression noise level equal to 10. The noise’s standard deviation decreases as the frequency increases, as suggested by our model. However, this effect is not homogeneous across different JPEG qualities, we can notice that the lower the JPEG quality, the more notorious this effect becomes. It is also remarkable that lowering the JPEG quality more strongly damages the high frequencies of noise, while keeping the low ones mostly unchanged. Furthermore, for a compression quality $QF = 90$ we observe that the effects of compression are pretty innocuous as the noise standard deviation is still mostly homogeneous across frequencies.

Figure 5 shows the standard deviation of the noise after JPEG compression divided by q as a function of the pre-compression noise divided by q , for different input images and quantization factors $q = 8, 64$. We observe that for very small values of noise, the output noise is bigger than the input noise on real images. Although the theory suggests output noises should be smaller in this range of values, the experimental results are disturbed by JPEG noise. This JPEG noise is negligible when compared to higher noise levels but becomes predominant when the noise levels are very low. For medium noise levels we observe, as suggested by our model, that output noises are smaller than input noises. This is due




Fig. 5: Ratio of the post-compression noise standard deviation and quality factor $\frac{\sigma}{q}$ as function of the same ratio before compression, with $\sigma \in [1, 20]$ and $q = 8$ (left), $q = 64$ (right). The model curve corresponds to Eq. 5 and highly coincides with the pure noise curve.

to the effect of quantization since it removes the small variations in pixel’s values for which noise is responsible. As the noise grows, we observe that the curves converge to the identity when q is low enough, as predicted by the model on pure noise. This is explained by the fact that, for big enough noise levels, the quantization factor becomes negligible with respect to the noise level. However, with a larger value of q , the noise seems dimmed by compression, deviating from the theoretical model. In practice, this is only caused by clipping; indeed to reach a $\frac{\sigma}{q}$ ratio of 1 when $q = 64$, the noise’s standard deviation must also reach 64. More pixels thus become saturated when noise is applied, thus lowering the noise level before even compression. When comparing the output noise to the post-clipping, pre-compression noise, we fall back to a curve similar to the $q = 8$ case. Although the two curves are similar, a given noise level will have a lower $\frac{\sigma}{q}$ ratio if q is high, and the noise will be more affected. Since high-frequencies components are often attributed higher quantization factors, their noise is thus reduced more.

On the other hand, we observe that textured images, such as traffic, show bigger deviations from the theoretical model than non-textured images or even pure synthetic noise. This phenomenon can be explained by the fact that in textured images high frequencies are not only affected by noise but also by texture. Even though our method removes most of the quantization noise, some remain in the residual, and is more prominent when the noise level is comparatively lower.

6. DISCUSSION

In this article, we derived a model for the effect of JPEG compression on prior noise. Both the theoretical and experimental results show that post-quantization, prior noise is frequency-dependent. In particular, previously-normal noise only remains normal separately for each DCT coefficient. We believe the applications for our study to be numerous and varied, in particular whenever precise knowledge of the noise of an image is required, such as restoration of bad quality (compressed) images, steganography, and image forensics. Our

work focused on the luminance channel. Future work will extend this model to the chroma components, which are sub-sampled in addition to the quantization we studied here.

References

- [1] A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian, "Practical poissonian-gaussian noise modeling and fitting for single-image raw-data," *IEEE transactions on image processing*, vol. 17, Nov. 2008.
- [2] N. N. Ponomarenko, V. V. Lukin, S. K. Abramov, K. O. Egiazarian, and J. T. Astola, "Blind evaluation of additive noise variance in textured images by nonlinear processing of block DCT coefficients," in *Image Processing: Algorithms and Systems II*, vol. 5014, SPIE, 2003.
- [3] N. N. Ponomarenko, V. V. Lukin, M. Zriakhov, A. Kaarna, and J. Astola, "An automatic approach to lossy compression of aviris images," *2007 IEEE International Geoscience and Remote Sensing Symposium*, 2007.
- [4] G. Yu and G. Sapiro, "DCT image denoising: a simple and effective image denoising algorithm," *IPOL*, vol. 1, 2011.
- [5] D. L. Donoho and I. M. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *JASA*, vol. 90, no. 432, 1995.
- [6] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising," *IEEE TIP*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [7] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-d transform-domain collaborative filtering," *IEEE TIP*, vol. 16, no. 8, pp. 2080–2095, 2007.
- [8] D. Cozzolino, G. Poggi, and L. Verdoliva, "Splicebuster: A new blind image splicing detector," Nov. 2015.
- [9] M. Gardella, P. Musé, J.-M. Morel, and M. Colom, "Forgery detection in digital images by multi-scale noise estimation," *Journal of Imaging*, vol. 7, no. 7, 2021.
- [10] T. Ehret, A. Davy, M. Delbracio, and J.-M. Morel, "How to Reduce Anomaly Detection in Images to Anomaly Detection in Noise," *IPOL*, vol. 9, 2019.
- [11] M. Chen, J. Fridrich, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE TIFS*, vol. 3, no. 1, 2008.
- [12] H. Gou, A. Swaminathan, and M. Wu, "Noise features for image tampering detection and steganalysis," in *2007 IEEE ICIP*, vol. 6, 2007.
- [13] B. K. T. Ho, V. Y. Tseng, M. Ma, and D. T. Chen, "Mathematical model to quantify JPEG block artifacts," in *Medical Imaging 1993: Image Capture, Formatting, and Display*, vol. 1897, SPIE, 1993, pp. 269–275.
- [14] B. Li, T.-T. Ng, X. Li, S. Tan, and J. Huang, "Revealing the trace of high-quality jpeg compression through quantization noise analysis," *IEEE TIFS*, vol. 10, no. 3, pp. 558–573, 2015.
- [15] —, "Statistical model of jpeg noises and its application in quantization step estimation," *IEEE TIP*, vol. 24, no. 5, 2015.
- [16] N. Le and F. Retraint, "An improved algorithm for digital image authentication and forgery localization using demosaicing artifacts," *IEEE Access*, vol. 7, 2019.
- [17] S. Mandelli, N. Bonettini, P. Bestagini, and S. Tubaro, "Training cnns in presence of jpeg compression: Multimedia forensics vs computer vision," in *2020 IEEE WIFS*, 2020.
- [18] W.-L. Lau, Z.-L. Li, and K.-K. Lam, "Effects of jpeg compression on image classification," *IJRS*, vol. 24, no. 7, pp. 1535–1544, 2003.
- [19] T. Jullian, V. Nozick, and H. Talbot, "Image noise and digital image forensics," in *Digital-Forensics and Watermarking*, Y.-Q. Shi, H. J. Kim, F. Pérez-González, and I. Echizen, Eds., 2016, pp. 3–17.
- [20] Y. Jiang, H. Zeng, X. Kang, and L. Liu, "The game of countering jpeg anti-forensics based on the noise level estimation," in *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, 2013.
- [21] S. Corchs, F. Gasparini, and R. Schettini, "Noisy images-JPEG compressed: subjective and objective image quality evaluation," in *Image Quality and System Performance XI*, vol. 9016, SPIE, 2014.
- [22] M. Delbracio, D. Kelly, M. S. Brown, and P. Milanfar, "Mobile computational photography: A tour," *Annual Review of Vision Science*, vol. 7, no. 1, 2021.
- [23] G. Fechner, "Elemente der psychophysik, breilkopf und härtel," *Leipzig: Breitkopf und Härtel*, 1860.
- [24] F. J. Anscombe, "The transformation of poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3/4, 1948, ISSN: 00063444.
- [25] B. Widrow and I. Kollár, *Quantization Noise: Round-off Error in Digital Computation, Signal Processing, Control, and Communications*. 2008.
- [26] M. Colom, *Noise-free test images dataset*. [Online]. Available: http://mcolom.info/pages/no_noise_images/.