Improving Re-sampling Detection by Adding Noise

Lakshmanan Nataraj, Anindya Sarkar, B.S. Manjunath[*]

Dept. of Electrical and Computer Engineering, Univ. of California, Santa Barbara, CA USA 93106

ABSTRACT

Current image re-sampling detectors can reliably detect re-sampling in JPEG images only up to a Quality Factor (QF) of 95 or higher. At lower QFs, periodic JPEG blocking artifacts interfere with periodic patterns of re-sampling. We add a controlled amount of noise to the image before the re-sampling detection step. Adding noise suppresses the JPEG artifacts while the periodic patterns due to re-sampling are partially retained. JPEG images of QF range 75-90 are considered. Gaussian/Uniform noise in the range of 28-24 dB is added to the image and the images thus formed are passed to the re-sampling detector. The detector outputs are averaged to get a final output from which re-sampling can be detected even at lower QFs.

We consider two re-sampling detectors - one proposed by Poposcu and Farid [1], which works well on uncompressed and mildly compressed JPEG images and the other by Gallagher [2], which is robust on JPEG images but can detect only scaled images. For multiple re-sampling operations (rotation, scaling, etc) we show that the order of re-sampling matters. If the final operation is up-scaling, it can still be detected even at very low QFs.

Keywords: Digital Image Forensics, Re-sampling detection, Noise addition, JPEG Denoising

1. INTRODUCTION

Digital Image Forensics is a recent field aimed at detecting manipulations in digital images. Various techniques have been proposed to expose tampering [9] due to re-sampling, region duplication, lighting artifacts and JPEG compression artifacts. While creating image forgeries, re-sampling is very common since parts of the image are scaled/rotated/stretched. Hence, re-sampling detection is important forensic analysis.

Several algorithms have been proposed to detect re-sampling. In most cases, it is assumed to be done using linear or cubic interpolation. In [1], Popescu and Farid discuss how re-sampling introduces specific statistical correlations and show that they can be automatically detected using an Expectation-Maximization (EM) algorithm. The algorithm estimates the periodic correlations (after first detecting whether such a correlation exists) among the interpolated pixels - the specific type of the correlation indicates the exact form of the re-sampling. However, the EM-based method is very susceptible to JPEG attacks - especially when the JPEG quality factor (QF) is 95 or lower. The reason is that the periodic JPEG blocking artifacts interfere with the periodic patterns introduced by re-sampling. For images that are scaled using linear/cubic interpolation, Gallagher [2] proposed an algorithm by analyzing the variance of the second difference of interpolated signals. Although this method can detect only up-scaling, it is very robust against JPEG and detection is possible even at very low QFs. (Downscaled images can be detected up to a certain extent but not as robustly as up-scaled images.) This method was further improved by Mahadien et al [5] to tackle other forms of re-sampling using a Radon transform and derivative filter based approach. In [6], M.Kirchner showed a simpler method to improve [1].

We consider only Popescu's and Gallagher's detectors in our experiments. Popescu's detector has been well studied and is considered as state of the art to detect re-sampling in uncompressed images. Although Gallagher's detector can detect only scaling, it is very robust against JPEG. Further, the JPEG patterns can be clearly seen in the output. In this paper, we focus on robust re-sampling detection for JPEG compressed images and use the above detectors for analysis. We add controlled amounts of noise to the image before the passing the image to the re-sampling detector. Although this approach is counter intuitive, adding noise to enhance detection has been studied in the form of Stochastic Resonance(SR) noise[12]-[14]. In [11], Peng et al show how SR noise enhances the detection of micro-calcifications in digital mammograms, which in turn indicates an early sign of breast cancer.

In [10], we have shown that adding Gaussian noise can be used to "denoise" JPEG images, in scenarios like detecting image re-sizing, where the visual quality of the image need not be taken into account. On adding controlled noise (by adjusting the noise power), we showed that the JPEG periodic patterns were suppressed while the periodic re-sampling patterns were partially retained. In this paper, we do further analysis with Gaussian and Uniform noise and consider various re-sampling operations. We show statistically that adding noise to a JPEG compressed image reduces the artifacts due to JPEG and the noise added image behaves similar to an uncompressed image. From our experiments, we also find that when a series of re-sampling operations are done on an uncompressed image and then saved in JPEG format, scaling can be detected as long as it is the final re-sampling operation. Finally, we show that adding Uniform noise yields superior performance in detecting re-sampling when compared to the state of the art detectors.

The rest of the paper is organized as follows. In Section 2 we review the two re-sampling detectors, their performance on JPEG images with and without adding noise. In Section 3, we analyze the effects of adding noise to a JPEG image and show how it suppresses the blocking artifacts. In Sec.4, we propose a solution to detect re-sampling in JPEG compressed images. The experiments and results are detailed in Sec. 5

2. DETECTING RE-SAMPLING

2.1 Popescu and Farid's method

In [1], Popescu and Farid proposed the first method to detect re-sampling in digital images. We will provide a brief introduction to their approach. The idea is that re-sampling introduces periodic correlations among pixels due to interpolation. To detect these correlations, they use a linear model in which each pixel is assumed to belong to two classes- a re-sampled class M_1 and a non re-sampled class M_2, each with equal probability. The conditional probability for a pixel belonging to M_1 class is assumed to be Gaussian while the conditional probability for a pixel belonging to class M_2 is assumed to be uniform. To simultaneously estimate a pixel's probability of being a linear combination with its neighboring pixels and the unknown weights of the combination, an Expectation Maximization (EM) algorithm is used. In the Expectation Step, the probability of a pixel belonging to class M_1 is calculated. This is used in the Maximization step to estimate the weights. The stopping condition is enforced when the difference in weights between two consecutive iterations is very small. At this stage, the matrix (same dimensions as image) of probability values obtained in the Expectation step for every pixel of the image is called the "Probability map (p-map) ". For a re-sampled image this p-map is periodic and peaks in the 2D Fourier spectrum of the p-map indicate re-sampling. In the p-map, a probability value close to 1 indicates that a pixel is re-sampled.

The reader is further referred to [1] for a detailed description. Fig. 1:(a)-(c) shows a non re-sampled image ,its p-map and its corresponding 2D DFT while Fig.1:(d)-(f) shows the same for a rotated and cropped image(6 deg, bi-linear interpolation). Both images were saved in TIFF format. Peaks can be clearly seen in the p-map's spectrum for the re-sampled image. For better visibility, some low frequencies of the spectrum are zeroed out.

2.2 Gallagher's method

In [2], Gallagher proposed a simple algorithm to detect up-scaled images that were interpolated using linear or cubic interpolation. (Downscaled images can also be detected but the robustness is higher for up-scaled images.) He showed that the variance of the second difference of the interpolated signal has the same periodicity as the sampling rate of the original signal. To detect interpolation, the second derivative was computed along each row of the image and then averaged to obtain a mean trace of the second difference which is proportional to the signal variance [2]. This trace is periodic for linear/cubic interpolated images and peaks are present in the Fourier spectrum of the trace. Fig. 2 shows a portion of the mean trace of the second derivative (a) and its 1D normalized DFT (c) for a non re-sampled image and the same for an up-scaled (1.5, bilinear interpolation) image (Fig. 2 (b),(d)). Strong periodicity can be seen in the mean trace of the up-scaled image which is reflected in its Fourier spectrum (DC term is zeroed out for better visualization).
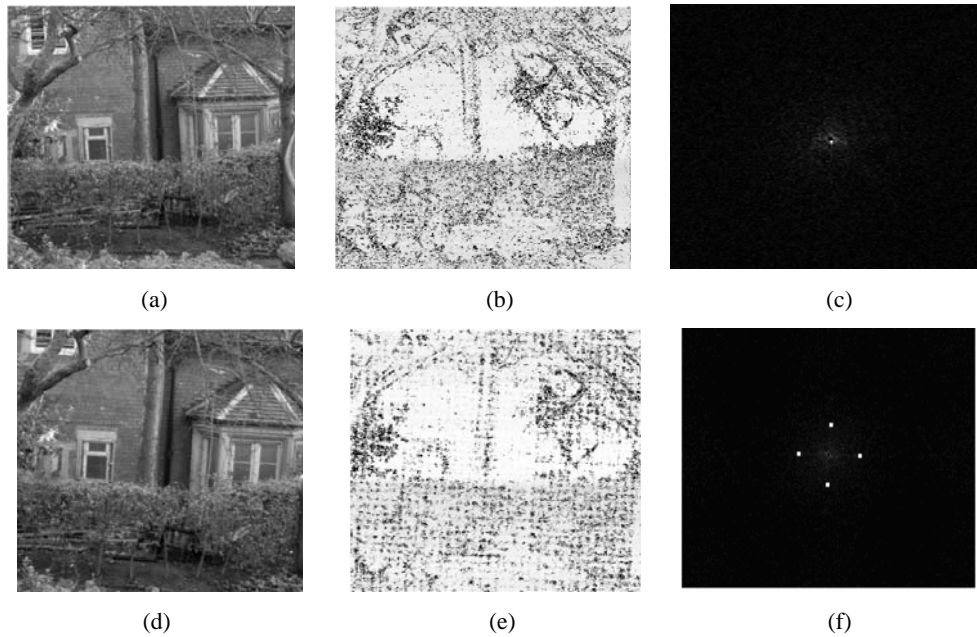
Fig. 1. (a)-(c): Original image, its p-map and its spectrum. (d)-(f): Image rotated by 6 deg, its p-map and its spectrum. Peaks can be seen in the p-map's spectrum for the rotated image.
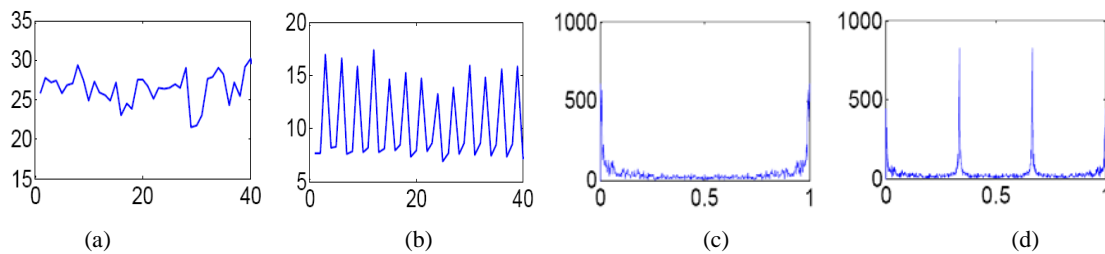

Fig. 2. (a),(b) Mean trace of second difference for original and up-scaled image. (1.5 bi-linear interpolation) (c),(d) Their corresponding normalized spectrum. Peaks can be seen in the spectrum for the up-scaled image.

2.3 Effect of Re-sampling Detection on JPEG images

As mentioned before, JPEG compression introduces blocking artifacts since it works on 8x8 blocks. This introduces periodic patterns which interfere with the periodic patterns introduced by re-sampling [1]. Hence, current re-sampling detectors are able to detect re-sampling only up to a JPEG QF of 95 [1],[5]. The periodic patterns introduced by JPEG can be seen clearly in the output of Gallagher's method (Fig. 3(a).) Peaks can be seen in the spectrum with a frequency spacing of 1/8. (Fig. 3(b)) In fact, this can also be used as a signature to detect previously compressed JPEG images, given a RAW bitmap!
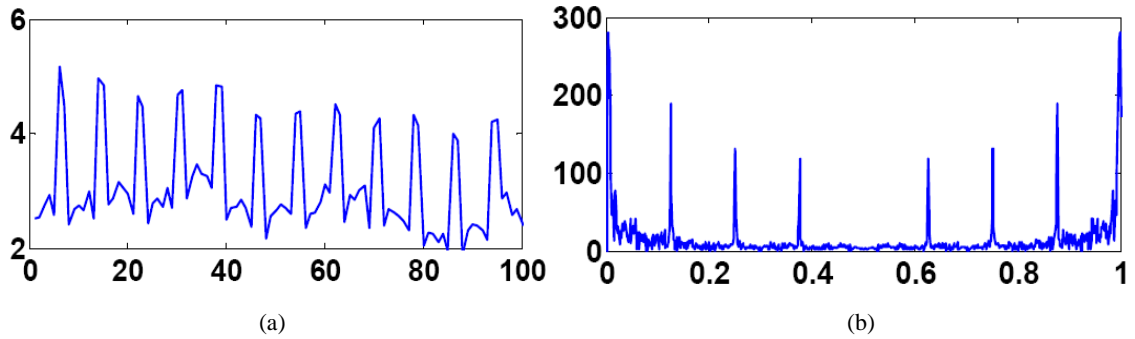
| (a) | (b) |

Fig. 3. (a) Mean trace of second difference for a non re-sampled JPEG compressed image (QF = 75), (b) its normalized spectrum

2.4 Adding Noise to suppress JPEG blockiness and detect re-sampling

We propose a simple method to suppress the JPEG artifacts by adding noise to the image pixels. While doing so, the periodic patterns due to re-sampling are not completely lost while the JPEG patterns almost disappear. However, while adding noise care is to be taken on the amount of noise that is added. If the Signal to Noise Ratio (SNR) is high, then it may not be enough to suppress the JPEG artifacts. If the SNR is too low, then it would also destroy the re-sampling patterns. From our experiments, we see that Gaussian/Uniform noise of SNR in the range of 24-28 dB suppresses the JPEG patterns for a moderately compressed JPEG image while partially retaining the re-sampling patterns.
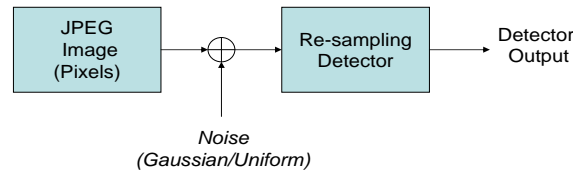

Fig. 4. Proposed method to detect re-sampled JPEG images

3.   ANALYSIS OF NOISE ADDITION

JPEG is a lossy compression scheme. If the original image is in an uncompressed format (say TIFF), information is lost when it is converted to JPEG. So now the question to be asked is what will happen to a JPEG image if we add noise-can we get back the original uncompressed image? Intuitively, this is not true since it is well known that quantization is irreversible. However, it could be possible that after adding noise the statistics of the image become much closer to the statistics of an uncompressed image since we see the JPEG artifacts disappearing. To analyze this further, we consider the problem of identifying whether a given bitmap image was previously JPEG compressed. The reason for considering this is two fold – First, the problem of identifying if a given bitmap was previously JPEG compressed or not has been well studied [3],[4]. Second, since we add noise in the pixel domain, the new matrix (image + noise) can be considered as a bitmap image. We follow the approach as mentioned in [3] for analysis and briefly review the method below.

3.1 Algorithm to detect JPEG compression history given a bitmap

Since JPEG works on 8x8 blocks, the algorithm to detect JPEG exploits the differences in pixel intensities within and across the blocks. To evaluate the algorithm, we use 800 uncompressed grayscale TIFF images of dimensions 384x384 and then save them in JPEG format at different QFs. The results shown below are averaged over the entire set. The input to their algorithm is already a bitmap in the case of uncompressed images, while for JPEG images they are converted from JPEG domain to pixel domain to form the bitmap. The algorithm is summarized as follows:

1. Divide the RAW bitmap image (matrix) into 8x8 blocks. (assuming block grid is known)

2. For every block, calculate the quantities

$$Z_1(i) \ = \ | \, A - B - C + D \, | \qquad Z_2(i) = \ | \, E - F - G + H \, | \qquad (1)$$

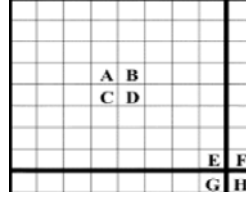where $i = 1 \ldots \#$ blocks and A-H are the pixel intensities as shown in Fig.5

Fig. 5. The square is an 8x8 block. **A-H** are the pixel intensities
for an 8x8 block and corresponding adjacent blocks.

3. The normalized histograms of the above quantities $H_I(n), H_{II}(n)$ are found and absolute histogram difference

$| H_I(n) - H_{II}(n) |$ is used to distinguish an uncompressed image from a JPEG. (Fig. 6(a))
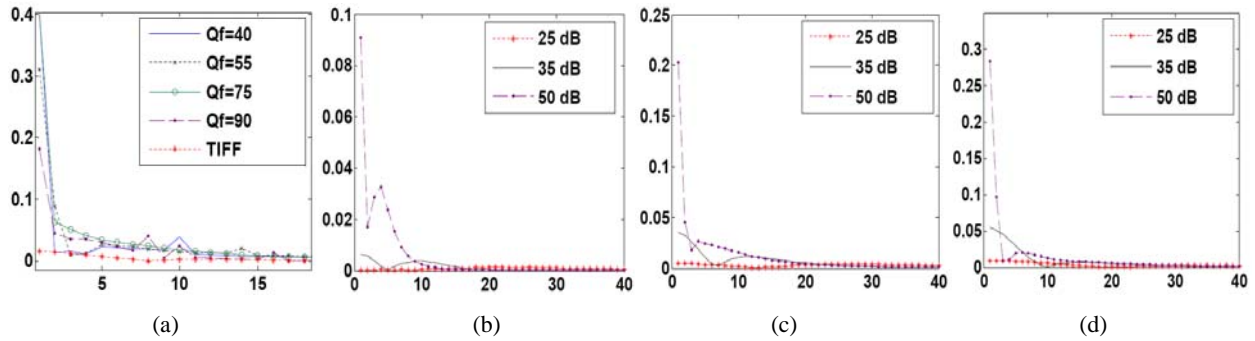

Fig. 6. Absolute histogram difference between: (a) TIFF and JPEG images at different QFs.
(b)-(d) JPEG images (QF 95,75 and 50) plus AWGN of 50,35 and 25 dB

As shown in Fig.6(a), if the image was previously JPEG compressed, it can be clearly identified using the above algorithm. The absolute histogram difference is almost zero for an uncompressed image because the pixel differences within a block and across corresponding adjacent blocks are similar. But for a JPEG image, the absolute histogram difference varies depending on the QF(Fig. 6(a)). To analyze the effect of noise addition to a JPEG image, we add Gaussian noise of varying strengths to JPEG images at various QFs in the pixel domain and repeat the same algorithm (Fig. 6 (b)-(d)). At a high SNR (50 dB), the noise added image is still similar to a JPEG image but as the SNR decreases we can see that it behaves more like an uncompressed image. However, this is not the original uncompressed image before JPEG compression or else JPEG would be a reversible process! In JPEG compression [7], information is lost when the DCT coefficients divided by the quantization matrix (of an 8x8 block) are quantized. More information is lost in the high frequency components since they are quantized to 0. Hence all 8x8 blocks in a JPEG image have a lot of zeros in the DCT domain which may contribute to the periodic blocking artifacts. Adding noise actually makes the coefficients non-zero which could possibly be the reason why the periodic artifacts are destroyed.

3.2 Effect of adding noise to a re-sampled and JPEG compressed image

In Sec. 3.1 we showed that adding noise to a JPEG image makes the noise added image similar to an uncompressed image. Now we would like to see what happens when noise is added to a re-sampled and JPEG compressed image. Consider an example where an uncompressed image is up-scaled by a factor of 1.4 using bi-linear interpolation. Now we consider three cases – first the scaled image is saved in TIFF format, then it is converted to JPEG at a QF of 75 and saved in JPEG format and finally Gaussian noise of 26 dB SNR is added to the JPEG image in the pixel domain. We then pass these three images through Gallagher's detector and examine their output in Fig. 7. On the left are the plots of the mean trace of the second difference and on the right are their corresponding spectrum. For the up-scaled uncompressed image, the periodic pattern can be clearly seen in Fig. 7(a). JPEG introduces its own periodicity which can be seen at normalized frequencies 1/8, 2/8 and 3/8 with 1/8 being the dominant frequency. (The peaks on the other half are symmetric since the image is real.) These interfere with the periodicity due to up-scaling of 1.4 as shown in Fig.7 (b,e) . On adding noise, the JPEG periodic artifacts are suppressed while the periodic patterns due to scaling are retained even though the magnitude of the peaks has dramatically decreased. However, adding too much noise may also suppress the periodic re-sampling patterns. The amount of noise that is to be added is discussed in the next section.
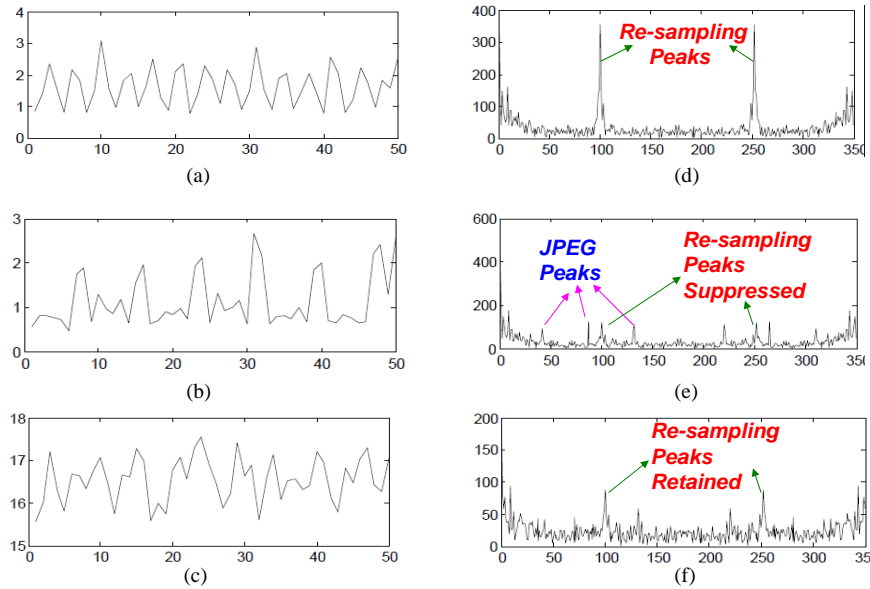

Fig. 7. Mean trace of second difference for: (a) Re-sampled image (b) Re-sampled and JPEG compressed image (c) Re-sampled + JPEG + noise. (d)-(f) Corresponding Spectrum.

3.3 Amount of Noise to be added

The amount of noise to be added depends on how severe the JPEG compression is. Adding too little noise (high SNR) will not be enough to suppress the JPEG artifacts. At the same time, if too much noise is added (low SNR), both the JPEG periodic artifacts as well as the periodic re-sampling patterns will be destroyed. We are interested in the range of noise that suppresses JPEG artifacts yet retains re-sampling patterns. From experimental results we found that for a moderate level of JPEG compression (QF 75-90), Gaussian/Uniform noise in the SNR range of 28-24 dB works well for different re-sampling operations and adding further noise destroys the re-sampling patterns. For lower QFs (QF 50-75) however, up-scaling can still be detected and the noise to be added can go up to an SNR of 20 dB. For example, consider

an image that is up-scaled to three times its original size using bi-linear interpolation and then JPEG compressed at QF=55. Fig. 8 displays the DFT of the mean trace of second difference at various levels of SNRs. At an SNR of 20 dB, most of the JPEG peaks (J1-J6) are suppressed while the re-sampling peaks (S1, S2) are retained. However at 15 dB SNR, even the re-sampling peaks are gone!
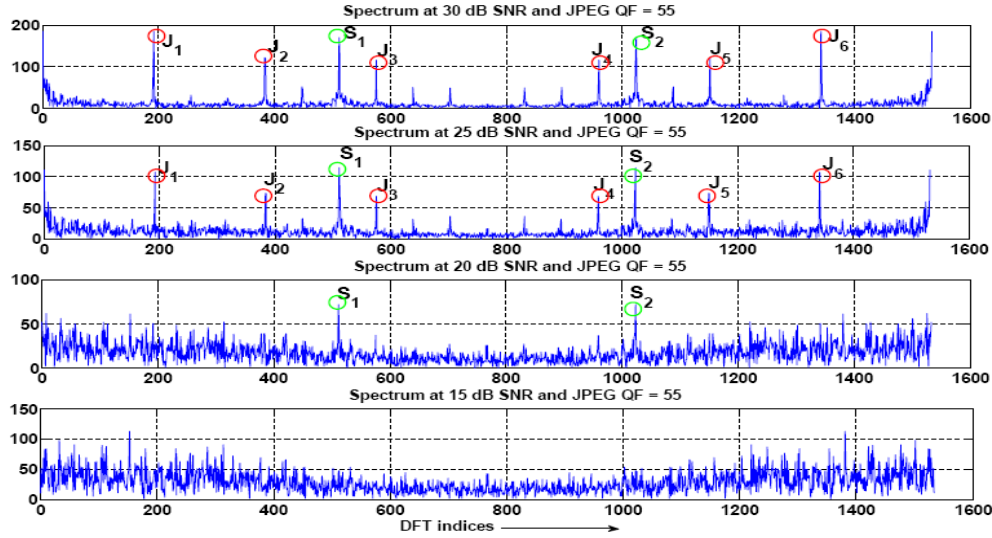

Fig. 8. By adding suitable amount of Gaussian noise, the peaks due to JPEG (J_1-J_6) are suppressed while still retaining the re-sampling peaks(S_1,S_2):-20 dB SNR. Below 20 dB, both JPEG and re-sampling peaks are suppressed.

4. PROPOSED METHOD

We propose a simple method to detect re-sampling in JPEG images. We add noise to the image in the pixel domain before passing it through the re-sampling detector. Since noise in the range of 28-24 dB works well for various re-sampling operations, we add noise in this range to generate 5 noise added images and then pass them through the detector. The detector outputs are averaged to get a final output (Fig.9). This averaged output is better than the output from a single detector. (shown later in Fig. 11)
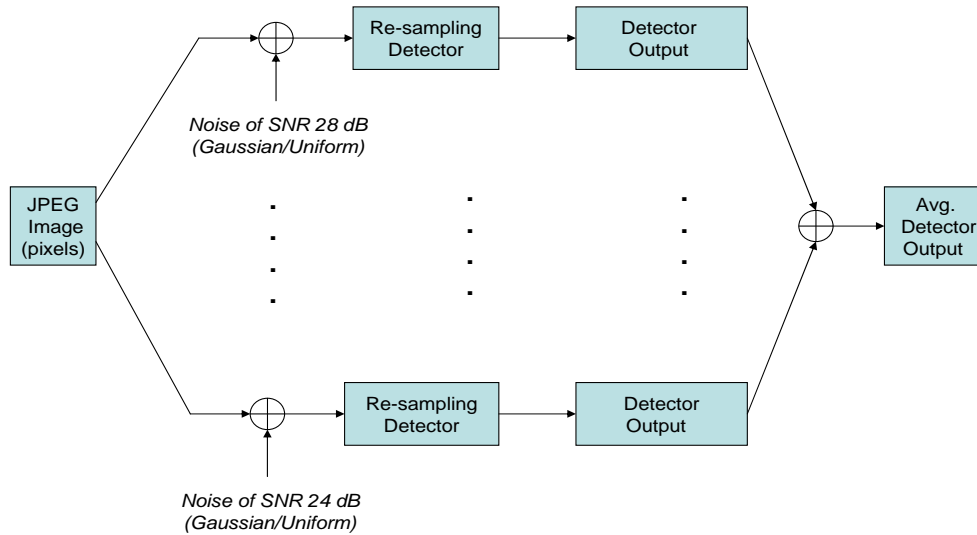

Fig. 9. Proposed approach to detect re-sampling in JPEG compressed images

5. EXPERIMENTS AND RESULTS

For our experiments, we use 800 uncompressed (TIFF) grayscale images of dimensions 384x384. The re-sampling operations done on the images are up-scaling, down-scaling, rotations, and affine transforms. The images are then saved in JPEG format at QF= 95, 90, 85, 80 and 75. We then add noise to the JPEG images in the pixel domain before passing it through the re-sampling detector as shown in Fig.9. Our results are summarized as follows.

- For scaling operations, we use Gallagher's re-sampling detector for evaluation since the scaling patterns can be clearly seen in its output even at low QFs (Fig. 8). From our experiments, we find that as long as the final operation is scaling before JPEG compression step, scaling can always be detected since it creates stronger patterns compared to other re-sampling operations- the only exceptions being the scaling factors where scaling peaks and JPEG peaks coincide. For example, consider a case where an uncompressed image is first up-scaled and then rotated (Fig. 10 (a)) and another case where the same image is first rotated and then up-scaled (Fig.10 (b)). Both the images are saved in TIFF format and visually, they appear almost the same. In the absence of JPEG, peaks can be seen at the detector's output for the second case (Fig. 10 (d)) because of the periodic patterns created due to up-scaling the rotated image , while the peaks are absent for the first case (Fig. 10 (c)). Now let the two images be JPEG compressed (QF = 75). When the two JPEG images are examined, visually they will still be similar but the re-sampling patterns will be present only for the second case (Fig.10 (e)). But now, even the JPEG peaks are present. After adding noise, the JPEG peaks are also suppressed (Fig. 10(f)).
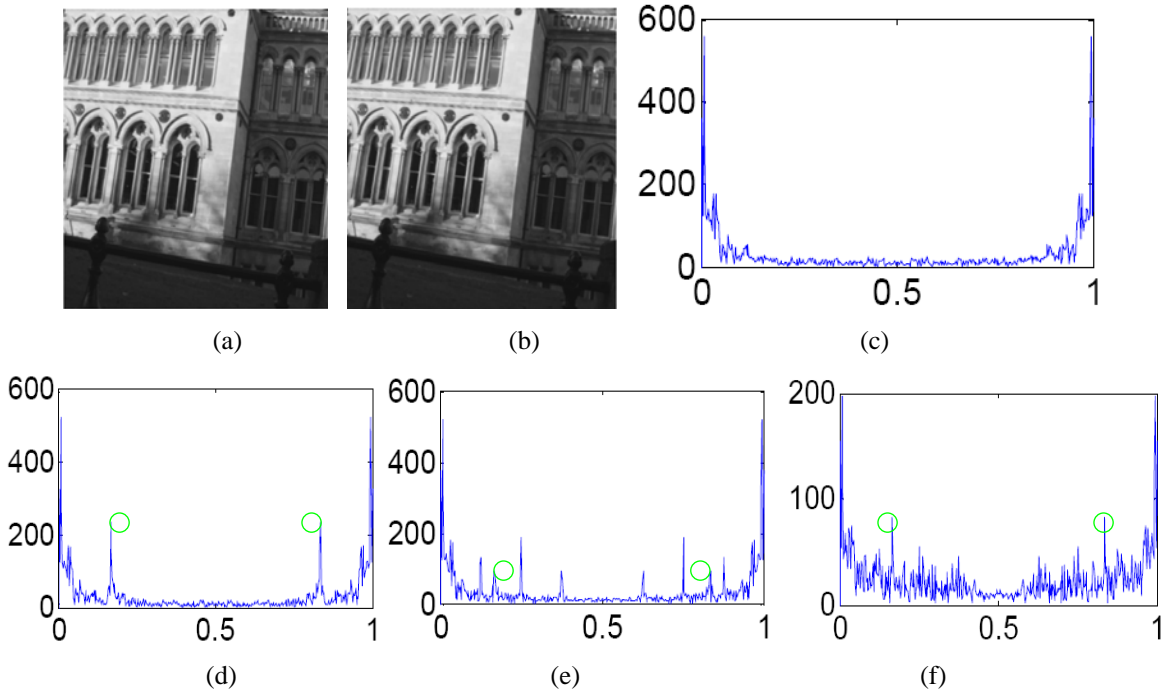

Fig. 10. (a) Image first scaled and then rotated. (b) Image first rotated and then scaled. (c-f) Re-sampling detector's output for first image, second image, after JPEG compressing the second image (QF=75) and after adding noise to the JPEG image. The JPEG peaks almost mask the scaling peaks but adding noise suppresses the JPEG peaks while retaining the scaling peaks. Green circles denote scaling peaks.

- For rotations and other affine transforms, we consider the spectrum of Popescu's p-map (Fig. 11(a)). For illustration, consider an example where an uncompressed image is rotated by 15 degrees and then JPEG

compressed at QF=80. JPEG causes two types of periodic patterns: strong periodic patterns which can be seen as high magnitude peaks in the horizontal and vertical axis of the p-map's spectrum and weak periodic patterns that are present all over the p-map's spectrum as peaks of smaller magnitudes at frequency spacing of 1/8 along horizontal and vertical directions(Fig. 11 (b)) (The strong pattern can also be seen in Gallagher's output at normalized frequency 1/8 while the weak patterns are the peaks of diminishing magnitude at frequencies 2/8 and 3/8: Fig.3 (b)). This makes re-sampling harder to detect since the peaks introduced by re-sampling are masked. When we add noise, most of the weak periodic patterns are suppressed and the re-sampling patterns slightly re-appear (Fig. 11(c)). However, the strong peaks present in the horizontal and vertical axis of the p-map's spectrum are not completely suppressed. Hence, for better visualization, the horizontal and vertical axes are zeroed out along with the DC. (For rotations and affine transforms, the peaks are usually not on the vertical/horizontal axes and one can also zero out only the JPEG peaks instead of zeroing out the complete axes.) Finally, five spectrums are found for the five noise added images (28-24 dB SNR) as shown in Fig.9 and the spectrums are averaged to get a final spectrum. (Fig. 11(d)). The re-sampling peaks can be seen more clearly in the averaged spectrum.
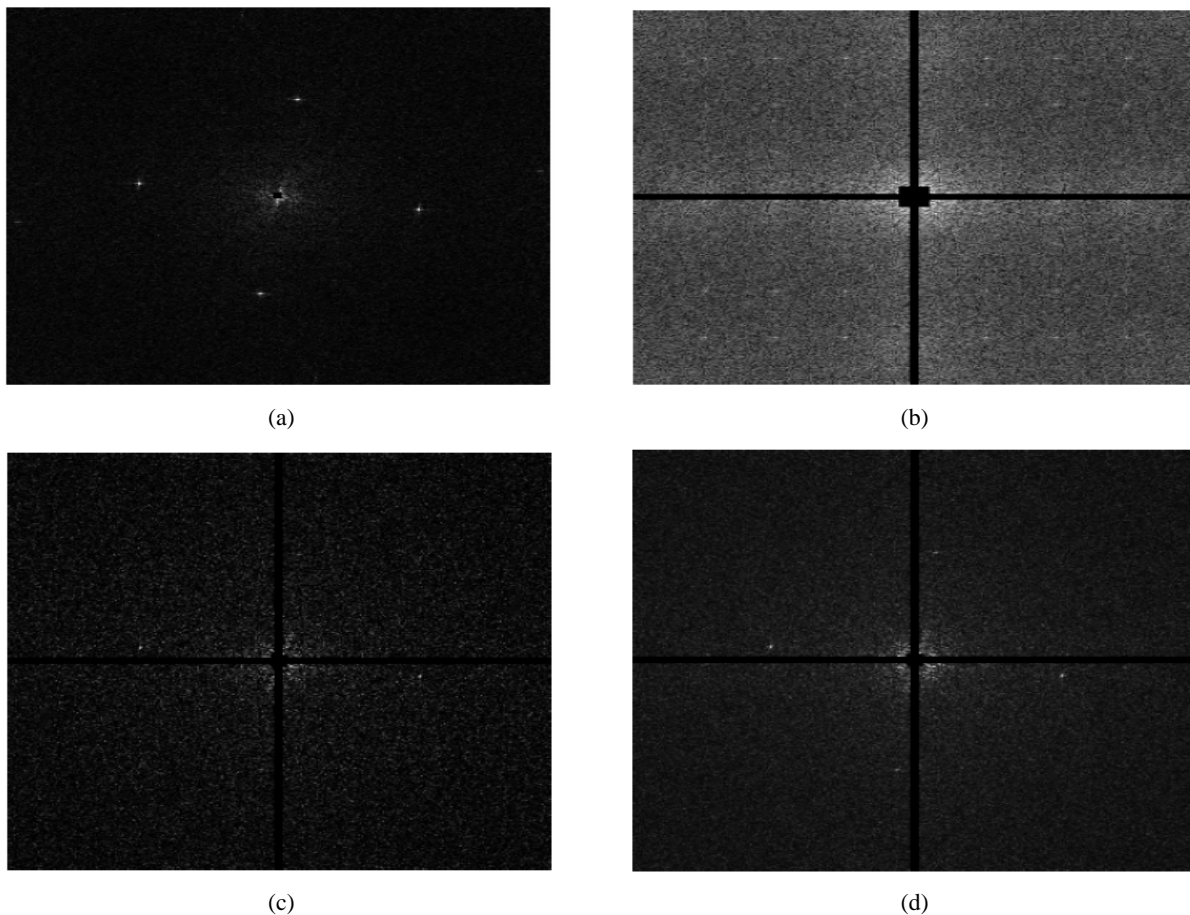

(a)   (b)

(c)   (d)

Fig. 11. (a) P-map's spectrum after rotation by 15 degrees. (b) P-map's spectrum after rotation followed by JPEG of QF=80. Multiple peaks can be seen all over showing weak periodic patterns of JPEG. (c) P-map's spectrum after rotation + JPEG + AWGN of 25 dB. Some rotation peaks are visible (d) P-map's spectrum is found for noise added images in the range of 28-24 dB SNR and then averaged .

We tested our method with various rotation angles at different JPEG QFs. The results were compared with Popescu and Farid's method, EM algorithm followed by JPEG de-blocking algorithm [8] and our method with Gaussian and Uniform Noise. We report the detection rates for a rotation angle of 25 degrees in Fig. 12 (a,b). We can see that our method of noise addition outperforms both Farid's original method and JPEG de-blocking method. Similar results were obtained for other rotation angles. For both bilinear and bicubic interpolation, addition of Uniform noise gave better results than Gaussian noise.
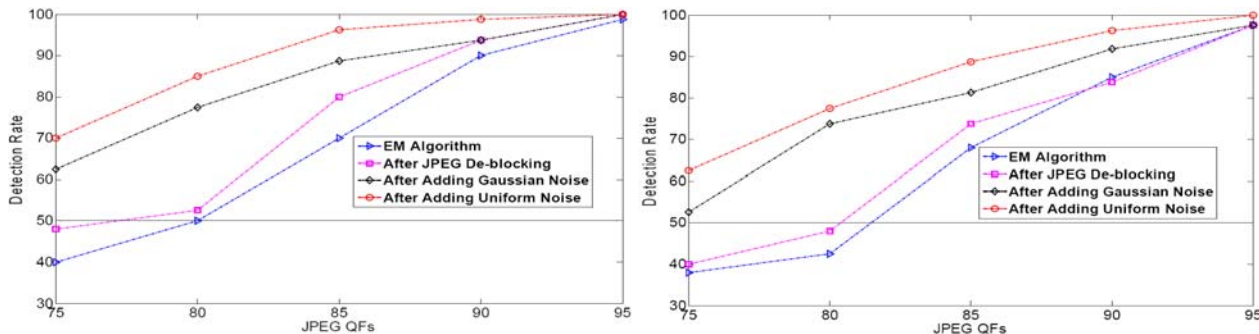

Fig. 12. (a,b) Detection rates for rotation angle of 25 degrees for bi-linear and bi-cubic interpolation.

6. CONCLUSION AND FUTURE WORK

We have shown that adding noise to a JPEG image in the pixel domain is an effective pre-processing step before passing the image through a re-sampling detector for forensic analysis. We show that the JPEG image after adding noise becomes similar to an uncompressed image. Addition of Uniform noise yielded superior results when compared with the original re-sampling detector. This makes re-sampling detection possible for moderately compressed JPEG images. Finally, we also find that scaling can be detected as long as it is the final re-sampling operation.

It can also be argued why is it even necessary to bother about JPEG images with JPEG2000 emerging as the standard in the near future. Further, JPEG2000 does not introduce as much distortion when compared to JPEG. However, from a forensic point of view, it is well known that current re-sampling detection methods fail on JPEG images. So even if JPEG2000 soon emerges as a standard, a forger can easily get away by creating a forgery, then running a JPEG compression algorithm on the image- knowing the re-sampling detection algorithms' weakness, and then convert it again to JPEG2000. To prevent scenarios like these, we feel that the problem of detecting re-sampling in JPEG compressed images should certainly be addressed.

In future, we will model the noise that is to be added and try to find the optimal pdf of the noise. It is well known that in JPEG images, most of the quantized block DCT coefficients are zeros. So the natural question that arises is can we do better if we add noise in the block DCT domain? We will analyze the effect of adding noise to those coefficients in the DCT domain and its effect on re-sampling detection.

REFERENCES

[1] A. C. Popescu and H. Farid, "Exposing Digital Forgeries by Detecting Traces of Re-sampling," IEEE Trans. on Signal Processing, vol. 53(2), 758-767, 2005.

[2] A. Gallagher, "Detection of linear and cubic interpolation in JPEG compressed images," Proc. Of Canadian Conf. on Computer and Robot Vision, 65-72, 2005.

[3] Zhigang Fan and R.L. de Queiroz, "Identification of bitmap compression history: JPEG detection and quantizer estimation," IEEE Trans. on Image Processing, vol. 12(2), 230-235, 2003.

[4] R. Neelamani, R.L. de Queiroz, Zhigang Fan, S. Dash. and R.G. Baraniuk, "JPEG compression history estimation for color images," IEEE Trans. on Image Processing, vol. 15(6), 1365-1378, 2006.

[5] B. Mahdian and S. Saic, "Blind Authentication Using Periodic Properties of Interpolation," IEEE Trans. on Information Forensics and Security vol 3(3), 529-538 ,2008.

[6] M. Kirchner, "Fast and reliable resampling detection by spectral analysis of fixed linear predictor residue," Proc. of the 10th ACM Workshop on Multimedia and Security, 11-20, 2008.

[7] G. Wallace, "The JPEG still picture compression standard," IEEE Trans. on Consumer Electronics, vol. 34(4), 30-44, 1991.

[8] A. Nosratinia, "Denoising of JPEG images by re-application of JPEG," Journal of VLSI Signal Processing, vol. 27, 69-79, 2001.

[9] H. Farid, "A Survey of Image Forgery Detection," IEEE Signal Processing Magazine vol 2(26), 16-25,2009.

[10] L.Nataraj, A.Sarkar, B.S.Manjunath, "Adding Gaussian Noise to "Denoise" JPEG images for Detecting Image Resizing," Proc. of ICIP, 2009.

[11] R.Peng, H.Chen, P.K.Varshney, "Noise-enhanced detection of micro-calcifications in digital mammograms, " IEEE J.Sel. Topics in Signal Processing, vol 3(1), 62-73,2009 .

[12] H.Chen, P.K.Varshney, S.M.Kay, J.H.Michels "Theory of the stochastic resonance effect in signal detection: Part I---Fixed detectors," IEEE Trans. on Signal Processing, vol 55, 3172-3184,2007 .

[13] H.Chen, P.K.Varshney "Theory of the stochastic resonance effect in signal detection: Part II---Variable detectors," IEEE Trans. on Signal Processing, vol 56, 5031-5041,2008 .

[14] S.Kay, "Can detectability be improved by adding noise?,", IEEE Signal Processing Letters, vol 7(1), 8-10, 2000.