

Robust Resampling Detection in Digital Images

Hieu Cuong Nguyen and Stefan Katzenbeisser

Computer Science Department, Darmstadt University of Technology, Germany
cuong@seceng.informatik.tu-darmstadt.de

Abstract. To create convincing forged images, manipulated images or parts of them are usually exposed to some geometric operations which require a resampling step. Therefore, detecting traces of resampling became an important approach in the field of image forensics. In this paper, we revisit existing techniques for resampling detection and design some targeted attacks in order to assess their reliability. We show that the combination of multiple resampling and hybrid median filtering works well for hiding traces of resampling. Moreover, we propose an improved technique for detecting resampling using image forensic tools. Experimental evaluations show that the proposed technique is good for resampling detection and more robust against some targeted attacks.

Keywords: Digital image forensics, resampling detection, targeted attack.

1 Introduction

With the availability of powerful tools for image processing, digital images can easily be altered without leaving visual evidence. Therefore, developing techniques for deciding on image authenticity became an urgent need. There are many different types of image tampering, which can be detected by different forensic methods. In order to create convincing forged images, manipulated images usually undergo geometric transformations, which require a resampling step. Thus, detecting traces of resampling became a popular approach in the field of image forensics. Techniques that detect resampling artifacts are often based on analyzing local linear dependencies [1, 2] or the variances of the second derivatives of images [3–5].

Robustness and security are important characteristics of any forensic detection technique along with detection capacity. While the robustness of a detection technique refers to the ability to detect forgeries even if the forged image is post-processed, the security refers to the ability to resist targeted attacks which were specifically tailored to disguise a forged image as authentic. In order to evaluate robustness, the aforementioned detection techniques were tested under several simple post-processing operations, such as Gaussian noise addition or JPEG compression. Recently, some authors [6, 7] designed testing frameworks to evaluate and compare the robustness of different resampling detectors. With the aim to assess the security of resampling detection, Kirchner and Boehme [8] designed several targeted attacks against the technique of Popescu and Farid [2]. Inspired by the above works, in this paper we propose some other simple but effective targeted attacks to conceal traces of resampling from common forensic tools.

Every existing resampling detector has its pros and cons. The technique of Popescu and Farid [2] is likely the most powerful but its use is complex and time consuming due to the use of the Expectation Maximization (EM) algorithm. In order to overcome the drawback of [2], Kirchner [1] proposed a fast detector which does not need to use the EM algorithm. Some other techniques based on detecting the variance of second derivatives in images are simpler to implement and provide faster detection in comparison with [2]. However, they suffer from high false positive rates and some of them [3, 4] are not capable of detecting rotated or skewed images. In this paper, we design an improved technique which is fast and robust in detecting resampled images. The technique is based on computing a so-called pseudo probability map of the image to be tested and applying the Radon transform to this map. The performance and security of the proposed technique are evaluated with a large image dataset under different attacks. Finally, we compare it with the state-of-the-art technique [2] under the same condition.

The structure of the paper is as follows. In the next section, we briefly review the concept of resampling and the main ideas of [2]. In Section 3, we propose some targeted attacks against resampling detection. After that, we present our improved resampling detection technique in Section 4. Experimental results will be shown in Section 5. Lastly, we conclude the paper in Section 6.

2 Techniques for Resampling Detection

2.1 Resampling and Interpolation

Once a geometric transformation such as scaling or rotation is applied to an image, a resampling process is involved. Interpolation is the central step of resampling in order to estimate the value of a signal at intermediate positions to the original samples. This step is the key to smooth the signal and then create a visually appealing image [9]. For example, a p/q resampling of an 1-D discretely-sampled signal consists of following three steps [2]:

- Upsampling: create a new signal $x_u[t]$, where $x_u[pt] = x[t]$, $t = 1, 2, \dots$ and $x_u[t] = 0$ otherwise.
- Interpolation: convolve $x_u[t]$ with $h[t]$: $x_i[t] = x_u[t] * h[t]$, where $h[t]$ is an interpolation filter (e.g. bilinear, bicubic).
- Downsampling: create a new signal $x_d[t]$, where $x_d[t] = x_i[qt]$, $t = 1, 2, \dots$

The extension to two dimensions is straightforward where the above mentioned operations are applied in both spatial directions.

2.2 Resampling Detection

There are several techniques to detect traces of resampling in digital images [1–5]. Among them, the technique of Popescu and Farid [2] is widely used and effective. The main step of [2] is to determine the probability of each sample being correlated to

its neighbors. To this end, the technique employs a linear predictor to approximate each sample's value as the weighted sum of its surrounding samples:

$$y_i = \sum_{k=-N}^N \alpha_k y_{i+k} + r_i . \quad (1)$$

The correlation probability p_i of each sample is computed based on the prediction error r_i , which is modeled as a zero-mean Gaussian random variable:

$$p_i = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-r_i^2}{2\sigma^2}\right) . \quad (2)$$

The probability values of all samples of an image together form the probability map (called p-map). The authors of [2] empirically found that the p-map of a resampled image is periodic and the periodicity becomes evident in the frequency domain by using the Fourier transform (DFT). However, the values of the weights (α) are usually not known in practice, so the p-map can not be computed directly. Therefore, the authors of [2] use an initial set of α for the estimation and then use Weighted Least Squares (WLS) integrated into an iterative EM algorithm in order to estimate the correlation of neighboring samples.

3 Attacks against Resampling Detection

The robustness of [2] was determined by applying different countermeasures, such as Gaussian noise addition and JPEG compression to resampled images. Nevertheless, Kirchner and Boehme [8] showed that the reliability of the technique was still solved only on the surface. Therefore, the authors proposed in [8] some targeted attacks against the technique [2]. The first attack is based on nonlinear filtering, the second attack is based on the Sobel edge detector, and the third attack integrates both mentioned attacks.

In this section, we design some other rather simple but effective targeted attacks against [2]. The first attack is based on multiple resampling by specific scales, the second attack is based on hybrid median filtering, and the third attack employs a combination of the attacks above. We also use the attacks to evaluate the security of our improved technique which we propose in Section 4. Experimental results will be presented in Section 5.

3.1 Attack Based on Multiple Resampling

When an image is downsampled by a factor of two, no sample in the downsampled image can be written as a linear combination of its neighbors [2]. Subsequently, traces of resampling should not be noticed in theory. Hence, we design an attack to disguise a resampled image by upsampling the image by a factor of two and downsampling it by a factor of two, thus yielding an image of the original size. In order to remove the aliasing artifacts of the downsampling process, the image is then anti-aliased.




Fig. 1. Shown in the top row is the original image, in the middle row the same image upsampled by a factor of 20%, and in the bottom row the same upsampled image, post-processed by the attack of multiple resampling. Each row shows the image itself, its p-map and the Fourier transform of the p-map.

Fig. 1 illustrates the detection process of [2] which consists of tested images, their corresponding p-maps and the Fourier transform of the p-maps. We realized that there is no peak in the Fourier transformed p-map of the original image, but in the case of an upsampled image, its transformed p-map has remarkable peaks. Although the quality of the tested image is not noticeably affected by the attack of multiple resampling, at the same time the peaks have not been absolutely eliminated (i.e. the traces of resampling can still be uncovered by the resampling detector). Using the detector of [2] on a dataset of 200 upsampled images by a factor of 20%, we obtained a detection rate of 99%. After applying the attack to the upsampled images, the detection rate is reduced to 84%.

3.2 Attack Based on Hybrid Median Filter

Since the technique [2] is based on detecting linear dependencies between samples in a locality, all kinds of nonlinear filters applied as a post-processing step are candidate attacks [8]. Kirchner and Boehme [8] proposed a targeted attack based on median filtering against [2]. While the attack is successful to conceal traces of resampling, the visual quality of attacked images suffers from noticeable blurring. To overcome this drawback, we design a targeted attack which based on another nonlinear filter called hybrid median filter [10]. The filter consists of three steps, each being applied to a $N \times N$ sliding window (N must be odd). In the first step one computes the median of

horizontal and vertical pixels in a $N \times N$ block (called M_1). In the second step we compute the median of diagonal pixels in the block (called M_2). Finally, the filtered pixel value is the median of the two median values (M_1 and M_2) and the center pixel of the block.

Fig. 2 illustrates the detection results of [2] for both kinds of nonlinear filters. We found that the median filter destroyed most evident peaks in the transformed p-map, but it also makes the image blurry. Conversely, the image attacked by the hybrid median filter is much less blurred, but sometimes peaks are still retained. When testing [2] on a dataset of 200 upsampled images by a factor of 20%, the detection rate is 99%. After applying the hybrid median filter to the upsampled images, the detection rate is degraded to 76%.




Fig. 2. Shown in the top row is the upsampled image attacked by the 3×3 median filter and in the bottom row the same upsampled image post-processed by a hybrid median filter with $N = 3$. Again, we show the image, its p-map and the Fourier transform of the p-map.




Fig. 3. Detection results of the upsampled image by a factor of 20% and then post-processed by the combination attack

3.3 Combination Attack

Although the proposed targeted attacks reduce the capability of detecting resampling, the detection rates are still high. In order to design a more powerful attack, we use them in combination: Firstly, the image is upsampled by a factor of two, then downsampled by a factor of two. The image is then anti-aliased. Lastly, a hybrid median filter is applied to the image.

Fig. 3 illustrates the detection results of an upsampled image which has been manipulated by the combination attack. We realized that all peaks disappeared in the transformed p-map, while the quality of the attacked image remains good. When we apply the combination attack to a dataset of 200 upsampled images by a factor of 20%, we found that the detection rate of the approach of [2] is reduced impressively to 3%.

4 An Improved Technique for Resampling Detection

4.1 Fast Resampling Detection

The core part of [2] is the EM algorithm used to estimate the probability of linear dependencies between neighboring samples. The results of all samples in the analyzed image are used to create the p-map. The remarkable peaks in the Fourier transformation of the p-map become evidence to uncover traces of resampling and can be recognized easily in the case of a resampled image.

Kirchner [1] showed that it does not matter what prediction weights (α) be used, the linear prediction errors which determine the p-map will be periodic in case of a resampled image. Thus, the author believed that the rather complex and time consuming EM estimation is not compulsory. As a result, he presented a fast but still reliable resampling detector.

Although the values of prediction weights (α) do not affect the periodicity of the p-map, different sets of α create different intensities in the p-map. For this reason, we call a p-map computed based on some pre-defined weights the pseudo p-map (pp-map for short). Through experiments, we found many times that using one predefined set of α for detecting an image by the technique [1], peaks can be recognized in the transformed pp-map, but using another set, peaks are not evident (though the periodicity exists in theory). Consequently, the selected set of α strongly affects the obtained outcomes. Whilst the major advantage of [1] versus [2] is bypassing the EM estimation, we believe that the technique [2], where the intensities of the p-map are correctly computed is more robust and reliable. Kirchner [1] empirically found one of the best preset filter coefficients α for computation of the prediction error as:

$$\alpha = \begin{bmatrix} -0.25 & 0.5 & -0.25 \\ 0.5 & 0 & 0.5 \\ -0.25 & 0.5 & -0.25 \end{bmatrix}. \quad (3)$$

4.2 Improved Resampling Detection

In this section, we introduce a resampling detection technique which consists of three main steps: computing the pp-map of the analyzed image, applying the Radon transform to the map and finding critical peaks in the transformed spectrum in order to infer the detection result.

Probability Map Computation. The residue of a sample is computed following Equation (1) where the weights (α) and the size of neighborhood (N) are pre-defined. The probability of correlation in a region $N \times N$ is estimated based on the residue, modeled as a zero-mean Gaussian noise described in Equation (2). These steps compute the pp-map (w) without using the EM algorithm as in [2]. The main steps of the algorithm are depicted in Algorithm 1, where $r(i)$ is the residue of a sample, $p(i)$ is the associated correlation probability and $w(i)$ is the corresponding in the pp-map.

Algorithm 1. Compute the pseudo probability map

Choose α , N , σ

Set $p_0 = 1/\text{max}_y$, where max_y is the size of the range of possible values for $y(i)$

for each sample i

$$r(i) = \left| y(i) - \sum_{k=-N}^N \alpha(k)y(i+k) \right|$$

$$p(i) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-r(i)^2}{2\sigma^2}\right)$$

$$w(i) = \frac{p(i)}{p(i) + p_0}$$

end

Radon Transformation. The Radon transform (RT) computes projections of an image along various directions given by a set of angles. The transformed result is the sum of the intensities of the pixels in each direction, i.e. a line integral [11]. The RT has robustness properties against rotation, scaling, and translation (RST) [12] and is also robust against additive noise [13].

Mahdian and Saic [5] improved the technique of Gallagher [4] by applying RT to the second derivatives of tested images. Accordingly, [5] can detect not only rescaled images but also rotated images. The major drawback of the technique [5] is its high false positive rate, especially in detecting images which contain strong textures. Inspired by the work of Mahdian and Saic [5], in our technique we apply RT to the pp-map of the image. To this end, firstly, the RT of the pp-map is computed for a set of predefined angles; this results in a set of projected vectors which are arranged in a matrix R . If the image has been resampled, the corresponding autocovariance matrix of the vectors contains a specific periodicity. Since our goal is to determine if an image has been subject to geometric transformations, we focus on the strongest periodic patterns present in the Fourier transform of the autocovariance of the projected vectors. Lastly, the strongest patterns are plotted in a spectrum from which the peaks are evident (see an example in Fig. 4 and Fig. 5). We assume that this technique works well for resampling detection due to the periodicity of the pp-map of resampled images shown in [1].




Fig. 4. Detection results of an original image. The peaks in the spectrum are not clear and distinguishable.

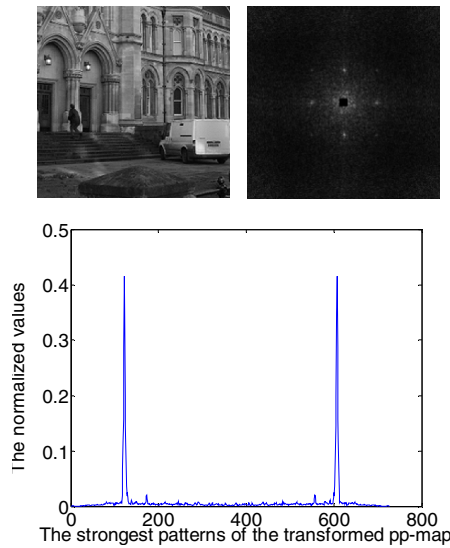


Fig. 5. Detection results of the upsampled image by a factor of 20%. The clear and strong peaks can easily be recognized.

Table 1. Detection rates when applying different attacks to upsampled images by a factor of 20%

	No Attack	Median Filter	Hybrid MF	Multiple Resampling	Combination Attack
[2]	99.0	1	76.0	84	3.0
Proposed	83.5	25	68.5	66	54.5

Peak Detection. After applying the RT (use the angles from 0° to 179° with an incremental step of 1°) to the pp-map, we obtain a spectrum where critical peaks can easily be recognized. If an image is resampled then there are clear and strong peaks in the spectrum. As an example, Fig. 4 and Fig. 5 show the results of applying the detector to an original image and a resampled image respectively. In order to infer the detection results, we search for strong peaks by computing the local maximums of the spectrum and choose the peaks based on a pre-defined threshold. The performance of the technique is improved when compared to [1, 2].

5 Experimental Results

In order to evaluate the detection techniques, we test them with different image datasets of original images, resampled images and attacked resampled images. Firstly, we randomly collected 200 uncompressed images from [14], converted them to gray-scale and cropped each of them to 256×256 pixels in order to create a dataset of original images. From the dataset of original images, we created different datasets of upsampled, downsampled, and rotated images by different factors (using bicubic interpolation).

In this section, we test our proposed technique and compare it to the technique of Popescu and Farid [2] as a baseline. We use the set of weights (α) as in (3) for the proposed technique. This set is also used as the initial weights in [2]. In both techniques, the size of the neighborhood is set to 3. In order to allow a fair comparison, we set their thresholds so that their detection rates in detecting upsampled images by a factor of 20% are larger than 80% and their false positive rates in detecting original images are lower than 5%.

As presented in Section 3, the median filter is a strong attack against resampling detectors based on measuring linear dependencies between neighboring samples. However, the major disadvantage of this attack is blurring. Among our targeted attacks, the hybrid median filter and multiple resampling affect image perception quality less, but they seem not strong enough. The combination attack is more powerful, while still maintaining the image quality. To confirm this, we apply the attacks to a set of 200 upsampled images by a factor of 20%. We test the attacked images with our proposed technique and the technique of [2]. The detection rates can be seen in Table 1. Both techniques work well to detect traces of resampling with detection rates of 99% and 83.5% respectively and the false positive rates below 5%. However, while the technique of [2] is mostly defeated by the combination attack with detection rate down to 3%, our proposed technique is much more robust, as the detection rate remains over 50%. Consequently, in this section, we use only the combination attack in order to evaluate the security of the resampling detection techniques.

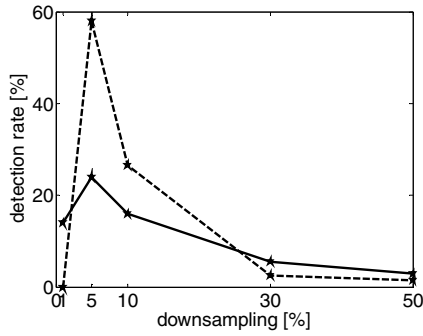


Fig. 6. Detection rates for downsampled images (dash line for [2] and solid line for the proposed technique)




Fig. 7. Detection rates for upsampled images (dash-star line for [2], solid-star line for the proposed technique) and for attacked upsampled images (dash-circle line for [2], solid-circle line for the proposed technique)




Fig. 8. Detection rates for rotated images (dash-star line for [2], solid-star line for the proposed technique) and for attacked rotated images (dash-circle line for [2], solid-circle line for the proposed technique)

Next, we test both techniques with downsampled images by different scaling factors. We realized that the detection rates of both techniques in detecting downsampled images are rather low (see Fig. 6). The reason is that the downsampling causes loss of information, thereby limiting the detection capacity of the statistical-based techniques.

We then evaluate the techniques with upsampled images and rotated images as well as their attacked versions. The attacked images are created by applying the combination attack to the resampled images. We found that both techniques can detect upsampled images by a scaling factor larger than 5% rather well (see Fig. 7). The technique of [2] even detects upsampled images by a factor larger than 10% perfectly (with a detection rate of nearly 100%). However, on the attacked images the detection rate of [2] is decreased significantly. This shows that [2] is not robust against this targeted attack. Although the proposed technique is not as powerful as [2] in detecting resampled images, it seems more robust against the combination attack. A similar situation occurs in detecting rotated images where both techniques work quite well in detecting rotated images by a factor larger than 3° (see Fig. 8). Although the proposed technique is little more robust than [2], both of them are almost defeated by the combination attack.

To assess the robustness of a detection technique, the authors usually test it with images where different post-processing operations have been applied. In this paper, we do not repeat the robustness evaluation of the original papers. However, we found an interesting property of the RT: its robustness against Gaussian noise is favorable for our technique. In other words, the proposed technique is less sensitive to noise. To confirm that, we test the techniques with upsampled images by a factor of 20% without any post-processing operation and with Gaussian noise addition. The results are shown in Table 2. While the detection rate of [2] is 99% in test with upsampled images, it is totally defeated when the images are post-processed by adding Gaussian noise by the Signal to Noise Ratio (SNR) of 20 dB.

A good attack not only reduces the detection rates of forensic techniques, but also maintains the image quality. There is usually a trade-off between the strength of attacks and the perceptual quality of the images which have been manipulated by the attacks. To quantify this aspect of an attack, we compute the average difference between pairs of resampled images (before the attack) and attacked resampled images (after the attack). The difference between a pair of images with the same size can be measured by calculating the PSNR (Peak Signal to Noise Ratio) or the Weighted PSNR (WPSNR). The WPSNR is an improved version of the PSNR firstly introduced in [15]. Based on the fact that the human eyes are less sensitive to modifications in textured areas than in smooth areas, the WPSNR uses an additional parameter called the Noise Visibility Function (NVF), which is a texture masking function. A higher PSNR or WPSNR usually indicates that the attacked image is of higher quality. In Table 3, we show the average PSNR and WPSNR of 200 pairs of upsampled images (by a factor of 20%) and their versions under different attacks of adding Gaussian noise (25 dB), median filtering and the combination attack. We found that the combination attack maintains the best image quality among the test cases.

Table 2. Detection rates for Gaussian noise added upsampled images by a factor of 20%

	No Attack	SNR 20 dB	SNR 25 dB	SNR 30 dB	SNR 35 dB
[2]	99.0	1.0	10	36	62.5
Proposed	83.5	36.5	68	77	79.0

Table 3. Average difference between resampled images and attacked resampled images (dB)

	Add Noise SNR 25 dB	Median Filter	Combination Attack
PSNR	21.20	20.29	22.93
WPSNR	34.30	32.74	36.13

6 Conclusion

In this paper, we revisited most important works for resampling detection in the literature. We designed some targeted attacks tailored to disguise traces of resampling in digital images. Since there is a relation between the derivative-based techniques and the techniques based on linear residue [1], we suppose that if the attacks can defeat [2], they will also work for attacking other resampling detection techniques. Subsequently, we proposed an improved resampling detection technique which consists of the steps of calculating the so-called pseudo p-map of the image, applying the Radon transformation and searching for critical peaks in the transformed spectrum. Since the proposed technique does not need the EM estimation to compute the pseudo p-map, it is much faster than [2]. We evaluated the performance and security of the proposed technique and the technique of Popescu and Farid [2]. We found that both techniques work well in absence of attacks and the technique [2] is the most powerful. However, our proposed techniques are more robust when attacks are applied.

References

1. Kirchner, M.: Fast and reliable resampling detection by spectral analysis of fixed linear predictor residue. In: Proceedings of the 10th ACM Workshop on Multimedia and Security - MM&Sec 2008 (2008)
2. Popescu, A.C., Farid, H.: Exposing digital forgeries by detecting traces of resampling. *IEEE Transactions on Signal Processing* 53, 758–767 (2005)
3. Prasad, S., Ramakrishnan, K.R.: On resampling detection and its application to detect image tampering. In: ICME (2006)
4. Gallagher, A.C.: Detection of Linear and Cubic Interpolation in JPEG Compressed Images. In: The 2nd Canadian Conference on Computer and Robot Vision (CRV 2005), pp. 65–72 (2005)
5. Mahdian, B., Saic, S.: Blind Authentication Using Periodic Properties of Interpolation. *IEEE Transactions on Information Forensics and Security* 3, 529–538 (2008)

6. Nguyen, H.C., Katzenbeisser, S.: Performance and robustness analysis for some resampling detection techniques in digital images. In: IWDW (2011)
7. Uccheddu, F., Rosa, A.D., Piva, A., Barni, M.: Detection of resampled images: performance analysis and practical challenges. EURASIP, 1675–1679 (2010)
8. Kirchner, M., Boehme, R.: Hiding Traces of Resampling in Digital Images. IEEE Transactions on Information Forensics and Security 3, 582–592 (2008)
9. Wolberg, G.: Digital Image Warping. IEEE Computer Society Press, Los Alamitos (1994)
10. Garcia, D.: BiomeCardio, <http://www.biomecardio.com/matlab/hmf.html>
11. Gonzalez, R., Woods, R., Eddins, S.: Digital image processing using Matlab. Gatesmark Publishing (2009)
12. Hoiland, C.: The Radon Transform. Aalborg University (2007)
13. Jafari-Khouzani, K., Soltanian-Zadeh, H.: Rotation-invariant multiresolution texture analysis using radon and wavelet transforms. IEEE Transactions on Image Processing 14, 783–795 (2005)
14. Schaefer, G., Stich, M.: UCID: an uncompressed color image database. In: Proc. SPIE, Storage and Retrieval Methods and Applications for Multimedia, San Jose, USA, pp. 472–480 (2004)
15. Voloshynovskiy, S., Herrigel, A., Baumgaertner, N., Pun, T.: A Stochastic Approach to Content Adaptive Digital Image Watermarking. In: Pfitzmann, A. (ed.) IH 1999. LNCS, vol. 1768, pp. 211–236. Springer, Heidelberg (2000)