

Counter-forensics of median filtering

D.T. Dang-Nguyen ^{#1}, I.D. Gebru ^{#2}, V. Conotter ^{#3}, G. Boato ^{#4}, F.G.B. De Natale ^{#5}

[#] *Department of Information Engineering and Computer Science, Univeristy of Trento
via Sommarive, 14 - 38123 Trento, Italy*

¹³⁴ {dangnguyen, conotter, boato}@disi.unitn.it

² israeldejene.gebru@studenti.unitn.it

⁵ denatale@ing.unitn.it

Abstract—Median filtering is a well-known non linear denoising filter often used as an harmless post-processing, sometimes also employed to affect the reliability of some forensic techniques. In this work, we present a novel counter-forensic method able to conceal the characteristic traces left by median filtering. By exploiting the knowledge of features used in existing median filtering detectors, we are able to remove the characteristic footprints via suitable random pixel modification, while keeping the quality of the counter-attacked image high. Experimental results show that the proposed method is very effective, computationally efficient and competitive with other state-of-the-art techniques.

I. INTRODUCTION

The pervasive availability of the Internet, coupled with the development of affordable and high-resolution digital cameras and sophisticated photo-editing tools, has lead to serious issues regarding the authenticity and fidelity of multimedia contents [1]. Indeed, the creation of visual compelling forged image has become an easy task even for a non-expert user, who can change the information they represent, without leaving any obvious traces of the occurred tampering.

In the context of digital multimedia security, digital image forensics operates, in contrast to active forensics techniques (i.e., digital watermarking), in absence of any special equipped device and not requiring the knowledge of any prior information about the content. The core assumption for this class of techniques is that original non-forged content owns some inherent statistical pattern introduced by the generative processing. Such patterns are always consistent in the un-forged content, but they are very likely to be altered after some tampering processes. Although visually imperceptible, such changes can be detected by means of an accurate statistical analysis of the content itself and taken as evidence of forgery.

Malicious manipulations of an image, such as splicing and region duplication operations (see for instance [2] and [3]), have been widely studied in the field of image forensics, since they alter the image content, both visually and semantically. However, even benign editing, such as median filtering [4][5][6], resampling [7] and compression [8], are of great interest from a forensic point of view. Indeed, even

though these operations are generally harmless and preserve the image content, they can seriously affect a forensic analysis of the content in various ways. First, it is in general very interesting to have the chance to study and reconstruct the history of the data (see, for instance, recent works supported by the European project REWIND¹). Second, similarly to what happens for steganalysis [9], the nature of an image prior to any manipulation (e.g., raw or compressed) may have a great impact on the performance of forensic algorithms. Finally, certain post-processing operations may alter or even erase the characteristic footprints left by manipulations and exploited by existing forensic detectors. As an example, in [10] the authors demonstrate how traces left by re-sampling operations can be made undetectable by post-processing the content with a median filter. In particular, median filtering is a very powerful tool in image processing, as it is a well-known denoising filter that preserve the image content and thus used to fool forensics techniques without affecting the image quality. Therefore, detection of median filtering can disclose possible counter-forensics operations. Indeed, in the last years, research in the new brand science of counter-forensics has attracted the attention of the scientific community. The final goal of counter-forensics is to identify weaknesses in existing forensic techniques, in order to assess and improve their trustworthiness [11]. Such studies may be of great help for researchers, pushing them towards improved techniques, able of both overcoming disclosed drawbacks and/or detecting when an anti-forensic tool has been used [12].

In this work, we propose a counter-forensic technique to conceal the characteristic traces left by median filtering and that can be used as a targeted attack against the state of the art median filtering detectors. Random pixel modification is introduced to remove the footprints that are searched by the available tools, while keeping a high fidelity of the post-processed image, compared to the median-filtered one. To the best of our knowledge, only one approach has been presented in the literature dealing with the hiding of median filtering in digital images [13]. The authors proposed an optimization problem aimed at designing a linear filter able to conceal fingerprints exploited by median filtering detectors while maximizing the quality of the counter-attacked image. With

extensive experimental results we prove that our proposed scheme outperforms the technique described in [13], both in terms of performances and computational complexity.

The basic idea of adding small perturbations interfering with the footprint left by filtering is pretty simple and makes the proposed algorithm computationally very efficient. Its application leads to a well performing tool, which we believe may represent a valuable contribution to the emerging field of counter-forensics for median filtering.

The structure of the paper is the following: in Section II we detail state of the art methods for median filtering detection, while in Section III we describe the existing counter-forensic technique for median filtering and propose our approach to hide traces of median filtering; in Section IV we report experimental results and finally in Section V we draw some conclusions.

II. MEDIAN FILTERING DETECTORS

Median filtering is a non linear operator widely employed for denoising and smoothing problems. Thanks to its non-linear nature, it is a valuable tool that can be used to limit the reliability of those forensic techniques based on some kind of linearity assumption [10]. Therefore, the blind forensics of median filtering is of particular interest and many detectors have been proposed in the literature.

To the best of our knowledge, three median filtering detectors working on uncompressed images have been proposed in the literature, namely [4], [5] and [6]. Following, we briefly review them, and later we will describe in details our proposed approach to defeat the efficiency of such algorithms. Recently, another technique has been proposed in [14], where authors propose to use median filtering residual for detection in JPEG compressed images. In this work we focus on uncompressed images but future work will be devoted to study counter-forensic technique in presence of JPEG compression.

A. Kirchner's method [4]

Kirchner et al. [4] exploit streaking artifacts as a characteristic fingerprint to detect median filtering in bitmap images. The basic idea is that median filtering introduces specific probability patterns (i.e., streaking artifacts), which yield to a non-zero probability that two output pixels with a certain distance originated from the same position of the input image. Such traces can be studied by means of the histogram \mathbf{h}_D of first-order differences. Authors demonstrate that the ratio between the bin centered in 0, h_D^0 , and the adjacent ones $h_D^{+1 \setminus -1}$ increases particularly in median filtered images. So, the ratio $\rho = h_D^0 / h_D^1$, is used as a discriminative statistic to distinguish between natural ($\rho = 1$) and filtered images ($\rho \gg 1$). However, since the classification based on ρ becomes unreliable for highly saturated original images, the authors propose to calculate the ratio ρ block-wise. They divide the image into a set of non-overlapping blocks of size $B \times B$ and determine the ratio ρ_b for each block. The final

discriminative measure $\hat{\rho}$ will be taken as the median of the weighted value obtained from all blocks:

$$\hat{\rho} = \text{median}_b(w_b \rho_b) \quad (1)$$

where the weights w_b are calculated as:

$$w_b = 1 - \frac{h_D^0}{B^2 - B}. \quad (2)$$

In this way, less weight is assigned to strongly saturated blocks, compensating for their effect.

B. Cao's method [5]

A similar basic idea was taken into account by Cao et al. in [5]. Authors exploit the fact that in original images the occurrence of equal neighboring pixels is more uncertain, especially in textured regions, while in median filtered images the difference between two adjacent pixels will be likely to be zero. Therefore, they take the probability of zero values on the first order difference map as a statistical feature to detect traces of median filtering. Given the image I , authors calculate, for each pixel (i, j) , the first-order row difference as:

$$\Delta I_r(i, j) = \begin{cases} 1 & \text{if } I(i+1, j) - I(i, j) = 0 \\ 0 & \text{if } I(i+1, j) - I(i, j) \neq 0 \end{cases} \quad (3)$$

Similarly, the first-order column difference ΔI_c is computed. To take into account highly textured regions in the image, a binary map $V(i, j)$ is computed for each pixel depending on the variance of a given neighboring region. The final scalar feature ρ is obtained as:

$$\rho = [f_r, f_c] \bullet [1/\sqrt{2}, 1/\sqrt{2}] \quad (4)$$

where

$$f_r = \frac{\sum_{i,j} \Delta I_r(i, j) \cdot V(i, j)}{\sum_{i,j} V(i, j)} \quad (5)$$

and f_c is calculated in the same way as f_r , but considering ΔI_c instead of ΔI_r .

C. Yuan's method [6]

The method presented in [6] is the most recent and elaborated one. It is based on the idea that median filtering, which is applied to overlapping blocks, affects the ordering of pixels within each block. Moreover, it introduces a strong dependence between median values originating from overlapping filter window. Such a dependence is characteristic of median filtered images and can be measured by means of 5 different sets of features:

- \mathbf{h}^{DBM} (Distribution of the Block Median): in median filtered images gray levels in a small block have a higher probability to be equal to the block median.
- \mathbf{h}^{OBC} (Occurrence of the Block Center Gray Level): gray level of the block center occurs more frequently in the block after filtering.
- \mathbf{h}^{QGL} (Quantity of Gray Levels in a Block): the quantity of gray levels in a block decreases after filtering given that the median filter does not produce new gray levels.

- \mathbf{h}^{DBC} (Distribution of Block Center Gray Level in Sorted Gray Levels): accounts for the distribution of block center gray level with respect to sorted gray levels (calculated as OBC but in the sorted gray levels).
- \mathbf{h}^{FBC} (First Occurrence of the Block Center Gray Level in Sorted Gray Levels): takes into account the first occurrence of the block center gray level in sorted gray levels.

Each subset of features captures the local pixel dependence introduced by median filtering. For a more detailed mathematical description of these feature, we refer the reader to [6].

Yuan proposes to fuse such features in order to create a scalar value f that can be used as an effective measure to discriminate between original and median filtered images:

$$f = \frac{h_5^{DBM} h_2^{OBC} h_6^{QGL} (h_3^{DBC} + h_7^{DBC} - h_2^{DBC} - h_8^{DBC}) h_3^{FBC}}{h_1^{OBC} h_9^{QGL} (h_2^{DBC} + h_8^{DBC} - h_1^{DBC} - h_9^{DBC}) h_2^{FBC} h_9^{FBC}} \quad (6)$$

where h_i^* is the i -th element of the array \mathbf{h}^* . For original non-filtered image the value of f is expected to be close to one, while the application of a median filter considerably increases its value. Such high-dimensional features based method provides a superior classification accuracy than single feature based approaches like [4] and [5].

III. COUNTER-FORENSIC OF MEDIAN FILTERING

Besides the described techniques, able to detect traces left by median filtering, very little work has been done so far on counter-forensics for hiding these traces. To the best of our knowledge, the only counter-forensic method has been proposed in [13]. It is a targeted, post-processing approach that attempts to modify the image so that the characteristic fingerprint left by median filtering is no longer detectable, while keeping the quality of the output image high. This is achieved through an optimization process, where the quality of the attacked image is maximized, still removing the traces of median filtering. Given the non-linear nature of the problem and the non-trivial issue of finding a good starting point for optimization, authors propose to use an iterative optimization algorithm to find a counter-attacked image, as follows:

- Select a set of median filtered images from the entire dataset (UCID database is used).
- Run an optimization process, looking for a linear processing operator to be applied to each image. Such operator should be able to interfere with the traces of median filtering, while maximizing the fidelity of the processed image. As a starting point for optimization, authors select a sharpening filter.
- Given the experiment where the best result is obtained (in terms of quality of the image and detectability of median filtering footprint), use the corresponding filter as a starting point for the optimization process over the rest of images.

Results obtained by the given approach are very satisfactory, even if we may argue that the entire process is computationally very expensive.

Since median filter has a block-wise application, most of the existing detection techniques are revealing the characteristic traces left by this filter using a block-wise analysis, instead of examining the whole image. For example, in [6], the presented method studies the behavior of pixels within each block, extracting a significant set of features, which are demonstrated to be discriminative between original and median filtered images. Therefore, counter-forensic methods aiming at hiding median filtering footprints would need to work block-wise as well. It would be beneficial to modify each block, in order to interfere with the traces left by previous median filtering and thus decrease the efficiency of the existing block-based detection methods.

A possible simple approach that can be used for this problem could be the addition of some random perturbation to the blocks. However, determining the proper amount of noise is not a trivial task: the addition of too much noise will unacceptably degrade the image quality, while adding too little noise could not be sufficient to hide the traces left by the median filter. As described above, in [13] the authors apply a linear filter on each block to minimize the effect of the features used by the detection methods and, in order to find a suitable filter kernel for each image, they exploit an optimization process. To overcome issues derived by this optimization process, we propose in this paper a simple but still effective counter-forensic method for median filtering which is computationally not expensive, outperforms performances of [13] and produces counter-attacked images with very high perceptual quality.

A. Proposed approach

The basic idea of the algorithm is the addition of a random noise into those blocks that are highly textured, so that degradation would be minimized, while still interfering with the characteristic footprint left by median filtering. Starting from the detector described in [6], we exploit the knowledge about the features \mathbf{h}^{OBC} (Occurrence of the Block center gray level) and \mathbf{h}^{DBM} (Distribution of the Block Median) to develop our counter-forensic technique. Since DBM features characterize the probability to have more than a single median values within a block, and OBC features take into account that the block center value can occur more frequently, these two sets of features are very representative of smooth blocks, where gray levels are likely to be equal to the block median. It would be beneficial not introducing any gray level modification in these blocks, since, even if we could be able to hide the traces of previous median filtering, unfortunately we will also introduce a visual degradation to the image (e.g., blocking artifacts and salt-and-pepper noise). In light of this, we discard those blocks which present high values (higher than 2) for elements of features \mathbf{h}^{OBC} and \mathbf{h}^{DBM} in (6), and consider the remaining set of blocks. Such group will be representative of highly textured areas in the image (e.g., edges) and we do add a random small perturbation to these blocks. This noise will still perturb characteristic features employed to detect median filtering ([4],[5] and [6]), but it will have a lower impact on

the quality of the counter-attacked image. Indeed, algorithm in [6] bases its detection on the discrimination feature f , as in (6), which is the sum of the contribution of different features extracted from each considered block. By adding more gray levels to those blocks which do not significantly contribute for OBC and DBM, we aim at minimizing the contribution due to the other three sets of features (especially QGL (Quantity of Gray Levels), but also DBC and DBC) to the final summation, thus decreasing the overall effectiveness of f . In the techniques presented in [4] and [5], the discrimination between original and median filtered images is based on the common idea that in median filtered images it is more likely to have equal neighboring pixels. Since our counter-forensic approach adds new gray levels to the selected blocks, pixel differences will be increased, that interfering with the discriminating features proposed in [4] and [5].

Unfortunately, the added perturbation may introduce brighter pixels in saturated blocks, thus compromising the quality of the final image. To work around this issue, we take into account the variance within a 3×3 neighborhood and decide to perturb only those blocks that present a standard deviation higher than an empirically fixed threshold T . Since the value of local standard deviation is higher for pixels on the edges, i.e., highly textured areas, introducing the dithering to those blocks will not have a strong impact on the visual quality of the image.

Shown in Algorithm 1 is the pseudo code of the proposed method. It takes a median filtered image I as an input and return a processed counter-attacked image I^c , which is a good approximation of I that would fool median filtering detection methods. The block size B and the threshold T are also required as the inputs. DBM and OBC features are computed as described in [6]. For each block b_k that does not contribute for these features, we calculate the standard deviation σ_{b_k} . If σ_{b_k} is higher than a certain fixed threshold T , we add a small perturbation to pixels in the block. We experimentally set the value of the added random noise η_1 to be in the range $[-7, -3] \cup [3, 7]$. A final loop (lines 13th-16th) is used to add some very small noise η_2 onto the whole image in order to further interfere with the median filtering traces, without affecting the final quality of the output image.

As reported in Algorithm 1, the main cost from a computation complexity point of view is the DBM and OBC computation which consists of $O(B^2)$ operations for each block. Hence the overall method has the complexity of $O(KB^2)$, where K is the number of the blocks in the image. Since only non-overlapping blocks are considered, $KB^2 \approx |I|$, where $|I|$ is the size of the image. Thus, the complexity of the proposed method is $O(|I|)$. It is worth noticing that the Fontani's method [13] requires a much higher complexity, due to the involved optimization process to build the filter. For each iteration it requires $O(|I|)$ operations, so in overall the method costs $O(\kappa|I|)$, where κ is usually much bigger than 1 and depends on the configuration of the Nelder-Mead Simplex Method applied for optimization (κ is the number of the searching points in it).

Algorithm 1 Proposed counter-forensic approach

Input:

- Median-filtered image $I(i, j) \in [0, 255]$
- Block size B
- Threshold T

Output: Counter-attacked image $I^c(i, j) \in [0, 255]$

Method:

```

1: initialize  $I^c$ :  $I^c \leftarrow I$ 
2: for each non-overlapping  $B \times B$  block  $b_k$  in  $I$  do
3:   if  $b_k$  does not contribute for DBM and OBC then
4:     compute standard deviation:  $\sigma_{b_k} \leftarrow STD(b_k)$ 
5:     if  $\sigma_{b_k} > T$  then
6:       for each pixel  $(i, j) \in b_k$  do
7:         randomly select  $\eta_1 \in [-7, -3] \cup [3, 7]$ 
8:         add noise to  $I^c$ :  $I^c(i, j) \leftarrow I(i, j) + \eta_1$ 
9:       end for
10:    end if
11:  end if
12: end for
13: for each pixel  $I^c(i, j)$  do
14:   randomly generate  $\eta_2 \in [0, 1]$ 
15:   add noise to  $I^c$ :  $I^c(i, j) \leftarrow I^c(i, j) + \eta_2$ 
16: end for

```

We believe that one of the biggest advantage of the proposed technique is its simplicity and low complexity. The basic idea of adding small perturbations in order to confound median filtering detectors may appear trivial, still its application leads to very effective results, as we are showing in Section IV. Moreover, research in the field of counter-forensic for median filtering is still in its early stage and this work may represent one of the first approaches contributing to the field.

IV. EXPERIMENTAL RESULTS

The main idea of the proposed forensic approach is to conceal traces of median filtering in uncompressed images. In order to verify the effectiveness of the proposed counter-forensic technique, we apply the three existing detectors for median filtering described in Section II to a dataset of images processed with a median filter and subsequently counter-attacked. The decreased accuracy of the detectors, together with the high perceptual quality of the counter-attacked image, will be the evidence of the effective action of the proposed method in hiding traces of median filtering.

We start considering all the 1338 uncompressed images present in the UCID dataset [15], which have also been used in the cited works about forensics and counter-forensics for median filtering. We perform a convolution with a median filter kernel of dimension 3×3 . Subsequently we apply the proposed counter-forensic technique, as described in Section III, and measure the detection accuracy of the forensic algorithms [4], [5] and [6] in terms of the resulting AUC (Area Under Curve) of the corresponding ROC plots. In Fig.1 we show the




Fig. 1. Accuracy performances of the three median filtering detectors (in terms of Area Under Curve), with respect to different settings of the threshold T .




Fig. 2. Average PSNR evaluated over the UCID database, with respect to different settings of the threshold T .

behavior of the three detectors with respect to different setting of the threshold T while fixing the block size $B = 3$. The horizontal axis corresponds to the threshold T , and the vertical axis to the classification accuracy of the three considered detectors. We clearly note that as the threshold T increases, we loose in performances, i.e., the detectors maintain reasonable discrimination rates.

Another important evaluation of the proposed method is the visual degradation introduced in the counter-attacked images. We measure the average PSNR (in dB) between the median-filtered and counter-attacked images and report results for different threshold T in Fig. 2. As expected, the PSNR monotonically increases with the threshold T . Similar results have been obtained employing a perceptual quality metric, i.e., Structural Similarity perceptual metric (SSIM) [16]. It is worth noticing, that similar behaviors have been tested when applying a median filter of size 5×5 . Moreover, we experimented that different settings for the block size B do not affect at all the general behavior of AUC and PSNR values reported in Fig.1 and 2, respectively.

Given the results on detection accuracy and average quality of the counter-attacked image, we selected $T = 2$ as the optimal setting to reach a good trade-off for our proposed algorithm in terms of ability to conceal traces of median filtering and fidelity of the modified image with respect to the




Fig. 3. ROC curves for the three median filtering detectors, evaluated on the UCID dataset, with threshold $T = 2$.

TABLE I
SHOWN ARE THE MEAN PSNR AND SSIM BETWEEN MEDIAN FILTERED AND COUNTER-ATTACKED IMAGES FOR THE THREE FORENSIC DETECTORS, EVALUATED ON THE UCID DATASET. ALSO SHOWN ARE THE ACCURACY PERFORMANCES, IN TERMS OF AUC, OVER MEDIAN FILTERED IMAGES (FIRST ROW) AND IMAGES COUNTER-ATTACKED WITH THE PROPOSED METHOD AND THE TECHNIQUES IN [13].

Method	PSNR	SSIM	AUC [4]	AUC [5]	AUC [6]
None	-	1	0.999	1.000	0.972
Proposed method	36.17	0.994	0.859	0.578	0.554
Sharpen [13]	23.3	0.815	0.923	0.525	0.706
Fontani [13]	30.77	0.940	0.924	0.679	0.709

median filtered one. Indeed, we are able to maintain an average PSNR of 36.17 dB and significantly decrease performances of the median filtering detectors. In particular, Fig. 3 shows the ROC curves for the three forensics methods [4], [5] and [6] evaluated on the counter-attacked images, when $T = 2$. For each method, we are able to decrease the accuracy to $AUC = 0.859$ for [4], to $AUC = 0.578$ for [5] and to $AUC = 0.554$ for [6], while keeping the average quality of the images high. It should be noted that the performances of the state-of-art forensic techniques evaluated over median filtered images are nearly optimal, as reported in Table I together with performances of the proposed method. From Table I we can state that the presented counter-forensic method is able to significantly affect the effectiveness of median filtering detectors, especially [5] and [6], while keeping a high visual quality of the counter-attacked image both in terms of PSNR and SSIM. Detector in [4] thus results to be the more robust one, as it had already been noticed when employing the anti-forensic tool in [13].

In order to further verify the efficiency of our method, we compare with the state-of-the-art counter-forensic technique proposed by Fontani in [13]. As reported in Table I, the proposed method outperforms Fontani's algorithm, both when it employs a simple sharpening filter or the optimized one. We perform better in terms of detection accuracy of methods [4], [5] and [6] evaluated on counter-attacked images. We especially reach excellent results in terms of average quality of



Fig. 4. In panel (a) is an example of median filtered image, while in panel (b) its counter-attacked version is reported. The small difference, mostly on edges, between the two images shown in panel (c) serves as demonstration of the high fidelity retained after the application of the proposed method.

the attacked image, outperforming Fontani's algorithm, with a gain in PSNR of almost 13dB compared to the Sharpening filter attack and a gain of almost 6dB with respect to the optimized algorithm in [13].

Finally, as an example, we show in Fig. 4 one median filtered image (panel (a)) and its corresponding version after the application of the proposed counter-forensic method (panel (b)). In panel (c), the logarithmic difference between the two images is shown, proving a high perceptual fidelity.

V. CONCLUSIONS

We have presented a novel counter-forensic method for hiding traces of median filtering. It is based on the idea that a random pixel modification can be introduced to perturb the characteristic footprints left by median filtering and that is exploited by existing median filtering detectors for uncompressed images. Experimental results show the effectiveness of the proposed approach, which is able to hinder performances of the state-of-the-art forensic detectors, while keeping the visual quality of the counter-attacked image high. We also demonstrated that our approach outperforms other available counter-forensic methods for concealing traces of median filtering. Based on its simplicity and effectiveness, the proposed approach represents a valuable contribution to the nascent field of counter-forensics for median filtering. Future work will be devoted to investigate counter-forensic techniques for JPEG compressed images.

ACKNOWLEDGMENT

The authors would like to thank Dr. Hai-Dong Yuan for providing the Matlab code of his method in [6].

REFERENCES

- [1] T. Gloe, M. Kirchner, A. Winkler, and R. Böhme, "Can we trust digital image forensics?" *15th International Conference on Multimedia*, vol. 0, pp. 78–86, 2007.
- [2] Y.-F. Hsu and S.-F. Chang, "Camera response functions for image forensics: An automatic algorithm for splicing detection," *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 4, pp. 816–825, 2010.
- [3] X. Pan and S. Lyu, "Region duplication detection using image feature matching," *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 4, pp. 857–867, 2010.
- [4] M. Kirchner and J. Fridrich, "On detection of median filtering in images," *Proc. of SPIE*, vol. 7541, pp. 101–1012, 2010.
- [5] G. Cao, Y. Zhao, R. Ni, L. Yu, and H. Tian, "Forensic detection of median filtering in digital images," *International Conference on Multimedia and Expo (ICME)*, pp. 89–94, 2010.
- [6] H. Yuan, "Blind forensics of median filtering in digital images," *IEEE Transaction on Information Forensics and Security*, vol. 6, no. 4, pp. 1335–1345, 2011.
- [7] H. C. Nguyen and S. Katzenbeisser, "Robust resampling detection in digital images," in *Communications and Multimedia Security*, ser. Lecture Notes in Computer Science, B. D. Decker and D. W. Chadwick, Eds., no. 7394. Springer-Verlag, 2012, pp. 3–15.
- [8] R. Neelamani, R. de Queiroz, Z. Fan, S. Dash, and R. G. Baraniuk, "Jpeg compression history estimation for color images," *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1365–1378, 2006.
- [9] R. B. A. Ker, "Revisiting weighted stego-image steganalysis," *Proc. of SPIE*, vol. 6819, pp. 501–517, 2008.
- [10] M. Kirchner and R. Böhme, "Hiding traces of resampling in digital images," *15th International Conference on Multimedia*, vol. 3, no. 4, pp. 582–592, 2008.
- [11] —, "Counter-forensics: Attacking image forensics," *Digital Image Forensics*, Springer, pp. 327–366, 2013.
- [12] M. Goljan, J. Fridrich, and M. Chen, "Sensor noise camera identification: Countering counter-forensics," *Proc. of SPIE*, vol. 7541, pp. 0S1–0S12, 2010.
- [13] M. Fontani and M. Barni, "Hiding traces of median filtering in digital images," *European Signal Processing Conference (EUSIPCO)*, pp. 1239–1243, 2012.
- [14] X. Kang, M. Stamm, A. Peng, and K. Liu, "Robust median filtering forensics based on the autoregressive model of median filtered residual," *Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pp. 1–9, 2012.
- [15] G. Schaefer and M. Stich, "Ucid - an uncompressed colour image database," *Proceedings of SPIE, in Storage and Retrieval Methods and Applications for Multimedia*, vol. 5307, pp. 472–480, 2004.
- [16] Z. Wang, A. Bovik, H. Sheikh, and E. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, 2004.