Hiding Traces of Resampling in Digital Images

Matthias Kirchner and Rainer Böhme

***Abstract*—Resampling detection has become a standard tool for forensic analyses of digital images. This paper presents new variants of image transformation operations which are undetectable by resampling detectors based on periodic variations in the residual signal of local linear predictors in the spatial domain. The effectiveness of the proposed method is supported with evidence from experiments on a large image database for various parameter settings. We benchmark detectability as well as the resulting image quality against conventional linear and bicubic interpolation and interpolation with a sinc kernel. These early findings on "counter-forensic" techniques put into question the reliability of known forensic tools against smart counterfeiters in general, and might serve as benchmarks and motivation for the development of much improved forensic techniques.**

***Index Terms*—Digital image forensics, resampling detection, tamper hiding, undetectable resampling.**

I. Introduction

OVER the past couple of years, digital imaging has matured to become the dominant technology for creating, processing, and storing the world's pictorial memory. This technology undoubtedly brings many advantages, but at the same time, it has never been so easy to manipulate images, often in such a perfection that forgery is visually indistinguishable from authentic photographs. As a result, the triumph of digital imaging harms the trustworthiness of pictures, particularly for situations in which society bases important decisions on them: in court (where photographs act as pieces of evidence), in science (where photographs provide empirical proofs), and at the ballot box (press photographs shape public opinion).

New streams of research have addressed the authenticity problem of digital images. A branch of it deals with tamper detection, which can be broadly subdivided into two main approaches. One approach is to track particularities of the image-acquisition process and report conspicuous deviations as indications for possible manipulation [1]–[4]. The second approach tries to identify traces from specific image-processing functions [5]–[8]. Although forensic toolboxes based on these approaches are already quite good at unveiling naive manipulations, we believe that they still solve the real problem only at its surface because little is known about the reliability of forensic techniques against a farsighted counterfeiter, who is aware of detection techniques.

In this paper, we change the perspective and introduce counter-forensic methods in the form of targeted attacks against a specific technique to detect traces of resampling in uncompressed images proposed by Popescu and Farid [6]. Section II contains a general consideration on the new subfield of research and its relation to more established disciplines as well as proposals for a harmonized terminology. Sections III and IV briefly recall the basics of interpolation methods and their detection before our countermeasures are presented in Section V. The design and results of a quantitative evaluation are reported in Section VI. Section VII discusses implications for future research on forensics and counterforensics.

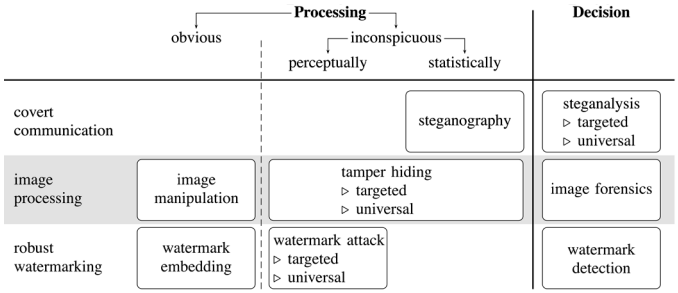

Fig. 1. Similarities and differences between tamper hiding, steganography, and attacks on robust digital watermarks.

II. Relations to Steganography, Steganalysis, and Digital Watermarking

Untraceable image manipulation, or *tamper hiding* [9], is a very young area of research. This justifies brief reflections on the relation to closely related fields in the area of multimedia security to develop consistent terminology (see Fig. 1).

Tamper hiding shares common goals with *steganography* [10]. Both try to achieve undetectability by preserving as many image properties as possible. Yet, steganography and tamper hiding differ in the amount and source of information to hide, and the extent to which an image can be altered. Most steganographic methods are designed to embed a given message by minimizing the number of changes to the cover (hence, keep its semantic) while tamper hiding conceals the mere information that larger parts of the original medium have been modified with the aim to change its semantic.

Steganalysis, as a counterpart to steganography, aims at unveiling the presence of a hidden message in a specific medium without having access to the original cover. A general analogy between *steganalysis* and *image forensics* becomes evident if we consider the act of forging images as information which is hidden inconspicuously in an image. Yet another parallel exists between tamper hiding techniques and attacks against *digital watermarking* schemes. Contrary to steganalysis, attacks against (robust) digital watermarks are designed to remove the embedded information rather than only detect it [11], [12]. In this sense, detectable manipulation artifacts can be understood

Manuscript received April 18, 2008; revised September 28, 2008. Current version published November 19, 2008. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Hany Farid.

The authors are with the Faculty of Computer Science, Technische Universität Dresden, Dresden 01062, Germany (e-mail: matthias.kirchner@inf.tu-dresden.de; rainer.boehme@inf.tu-dresden.de).

Digital Object Identifier 10.1109/TIFS.2008.2008214

as an inherent watermark, which tamper hiding techniques aim to remove or suppress. Therefore, we deem it appropriate to refer to specific tamper hiding methods as attacks against specific image forensic techniques.

Different tamper hiding techniques can be classified by their position in the process chain. We call a method *integrated* if it replaces or interacts with the image manipulation operation (e.g., an undetectable copy-move tool as a plug-in to image processing software) as opposed to *postprocessing*, which refers to algorithms that try to cover all traces after manipulation with conventional methods. Note that this distinction reflects the two different points of view on tamper hiding. An integrated attack is more closely related to steganography (hiding while processing) whereas postprocessing approaches resemble watermarking attacks (removing traces).

A second classification is borrowed from the context of steganalyis and watermarking attacks. We call a method targeted, if it avoids traces detectable with one particular forensic technique, which the developer of the attack usually knows. Conversely, universal attacks try to maintain or correct as many statistical properties of the image in order to conceal manipulations even when presented to unknown forensic tools. This appears to be the more difficult task: correct or at least plausible image statistics imply compliance with stochastic image models, which are not fully understood. So attackers can never be sure that a forgery is free of detectable artifacts.

Further interesting parallels are as follows.

- The design space for a steganographic algorithm is governed by a tradeoff between capacity and security. A similar relation can be found for attacks against digital image forensics. The stronger a manipulating operation interferes with the inherent image structure, the harder it is to feign an authentic image.
- Attacks against digital watermarks are often evaluated by the retained image quality after successful removal of a watermark [13]. Similarly, there appears to be a practical tradeoff between security (i.e., undetectability) and quality for tamper-hiding techniques. Plausible postprocessing, such as JPEG compression or blurring, may suppress detectable manipulation artifacts but, at the same time, go along with a loss in image quality. We conclude that every evaluation of attacks against digital image forensics should always be benchmarked against two criteria, namely, 1) undetectability and 2) image quality.

In the light of these considerations, we will now present a targeted tamper-hiding technique for undetectable resampling.

III. Preliminaries

Let $\mathbf{x} = (x_0, x_1, \ldots, x_{m_x n_x - 1})$, $x_i \in [0, 2^\ell - 1]$ be the column vector of integer pixel intensities of a $m_x \times n_x$ image, where the ith element holds the pixel value with coordinates $\mathbf{c} = (c_0, c_1)$, $0 \leq c_0 < m_x$, and $0 \leq c_1 < n_x$. Spatial coordinates are mapped to vector indices by a family of functions $\nu_m : \mathbb{N}^2 \to \mathbb{N}$, where m is the number of rows in the image

$$\nu_m(\mathbf{c}) = \mathbf{c} \cdot (1, m) = c_0 + c_1 m. \tag{1}$$

The inverse mapping is then defined as

$$\nu_m^{-1}(i) = \left(i - m \left\lfloor \frac{i}{m} \right\rfloor, \left\lfloor \frac{i}{m} \right\rfloor \right). \tag{2}$$

We write $\mathbf{y} = (y_0, y_1, \ldots, y_{m_y n_y - 1})$ as a $m_y \times n_y$ resampled version of a source image $\mathbf{x}$

$$\mathbf{y} = \mathbf{R}\mathbf{x}. \tag{3}$$

Matrix $\mathbf{R}$ has dimension $\dim(\mathbf{y}) \times \dim(\mathbf{x})$ and contains the weights of an interpolation filter $\phi : \mathbb{R}^2 \to \mathbb{R}$ with respect to an affine 2×2 transformation matrix $\mathbf{A}$. More specifically, element $R_{i,j}$ corresponds to the interpolation weight of original pixel x_j in the computation of transformed pixel y_i. Let $\delta_{\mathbf{A}, m_x, m_y} : \mathbb{N}^2 \to \mathbb{R}^2$ be a family of functions which return the relative position of an original pixel x_j and a transformed pixel y_i with respect to a transformation $\mathbf{A}$. For notational convenience, we omit indices m_x and m_y (number of rows in the original and transformed image, respectively)

$$\delta_{\mathbf{A}}(i,j) = \left(\mathbf{A} \cdot \nu_{m_x}^{-1}(j) - \nu_{m_y}^{-1}(i) \right) \odot \left(\operatorname{diag}(\mathbf{A}\mathbf{A}') \right)^{-1/2}. \tag{4}$$

Operator $\odot$ denotes element-wise multiplication and $(\operatorname{diag}(\mathbf{A}\mathbf{A}'))^{-1/2}$ is a normalization vector. The resampling matrix $\mathbf{R}$ is then specified by

$$R_{i,j} = \phi\left(\delta_{\mathbf{A}}(i,j) \right). \tag{5}$$

We use scaling and rotation matrices $\mathbf{A}^S$ and $\mathbf{A}^\Theta$ for upscaling ($S > 1$), downscaling without prefiltering ($S < 1$), and rotation by the counterclockwise angle Θ, $0 < \Theta \leq \pi/2$, respectively

$$\mathbf{A}^S = \begin{bmatrix} S & 0 \\ 0 & S \end{bmatrix} \quad \mathbf{A}^\Theta = \begin{bmatrix} \cos\Theta & -\sin\Theta \\ \sin\Theta & \cos\Theta \end{bmatrix}. \tag{6}$$

Unless otherwise stated, we use a linear interpolation kernel and set $\delta_{\mathbf{A}}(0,0) = (0,0)$.

IV. Detecting Traces of Resampling

Many attempts of image forgery rely on scaling and rotation operations, which involve a resampling process. As a result, scholars in image forensics have developed methods to detect traces of resampling in bitmap images. Present detectors rely on resampling artifacts observable in either the transformed image's derivatives [14]–[16] or in the residue of a local linear predictor [6]. Since the derivative-based approaches are not capable of detecting arbitrary affine transformations [14], [15], or suffer from high false positive rates [16], this paper focuses on Popescu and Farid's state-of-the-art detector [6]. It is known as a reliable and extensively tested [17] tool.

Interpolation algorithms are key to smooth and visually appealing image transformation [18]; however, a virtually unavoidable side effect of interpolation is that it creates linear dependences between adjacent pixels. As shown in [14] and [16], the strength of the linear dependence varies periodically with the cycle length, which itself depends on the resampling parameters. Popescu and Farid's detection method supports identifying the presence of such periodic artifacts. The intensity

of each pixel y_i can be modelled as a weighted sum of pixels in its $K \times K$ local spatial neighborhood (with $K = 2L + 1$ and L integer) plus a residual ε_i

$$y_i = \mathbf{P}^{\boldsymbol{\alpha},i} \cdot \mathbf{y} + \varepsilon_i. \tag{7}$$

The local linear predictor for pixel y_i, $\mathbf{P}^{\boldsymbol{\alpha},i}$ is defined as

$$\mathbf{P}^{\boldsymbol{\alpha},i} = \mathbf{1}^{1\times K^2} \cdot \left(\left(\mathbf{1}^{1\times \dim(\mathbf{y})} \otimes \boldsymbol{\alpha}\right) \odot N^i\right) \tag{8}$$

where $\mathbf{1}^{i\times j}$ is a $i \times j$ matrix of ones, and $\otimes$ denotes the Kronecker product. Vector $\boldsymbol{\alpha}$ contains the K^2 unobservable weights of $\mathbf{P}^{\boldsymbol{\alpha},i}$, with the center element $\alpha_{\lfloor K^2/2\rfloor} := 0$. We call $\mathbf{N}^i$ the $K^2 \times \dim(\mathbf{y})$ neighborhood matrix of pixel y_i, where the (k,j)th element $N^i_{k,j}$ is set to 1 if pixel y_j is the kth local spatial neighbor $k \in [0,\ldots,K^2-1]$ of pixel y_i

$$N^i_{k,j} = \begin{cases} 1, & \text{for } j = \nu_{m_y}\left(\nu^{-1}_{m_y}(i) + \nu^{-1}_{L}(k) - \nu^{-1}_{L}(\lfloor \frac{K^2}{2} \rfloor)\right) \\ & \quad \wedge k \neq \lfloor \frac{K^2}{2} \rfloor \\ 0, & \text{otherwise.} \end{cases} \tag{9}$$

A so-called p-map $\mathbf{p} \in [0,1]^{(m_x \cdot n_x)}$ is defined as a vector of probability measures for the strength of linear dependence for each pixel based on a simplified two-state model. It can be obtained from any given image as follows: Pixels y_i are assumed to belong to one of two sets $\mathcal{M}_1$ and $\mathcal{M}_2$. Set $\mathcal{M}_1$ contains all pixels with high linear dependence whereas set $\mathcal{M}_2$ comprises all pixels with low linear dependence. Popescu and Farid propose the expectation maximization (EM) algorithm [19], an iterative two-stage procedure, to estimate the probabilities for each pixel's assignment to $\mathcal{M}_1$, respectively, $\mathcal{M}_2$ and the unknown weights $\boldsymbol{\alpha}$. First, the E-step uses the Bayes theorem to calculate the probability for each pixel belonging to set $\mathcal{M}_1$

$$\begin{aligned} p_i &= \text{Prob}(y_i \in \mathcal{M}_1 \,|\, y_i) \\ &= \frac{\text{Prob}(y_i \,|\, y_i \in \mathcal{M}_1)\,\text{Prob}(y_i \in \mathcal{M}_1)}{\sum_{k=1}^{2} \text{Prob}(y_i \,|\, y_i \in \mathcal{M}_k)\,\text{Prob}(y_i \in \mathcal{M}_k)}. \end{aligned} \tag{10}$$

Evaluating this expression requires:

1) a conditional distribution assumption for $\mathbf{y}$: $y_i \sim \mathcal{N}(\mathbf{P}^{\boldsymbol{\alpha},i} \cdot \mathbf{y}, \sigma_{\mathcal{M}_1})$ for $y_i \in \mathcal{M}_1$ and $y_i \sim \mathcal{U}(0, 2^\ell - 1)$ for $y_i \in \mathcal{M}_2$, where $\mathcal{N}(\mu,\sigma)$ denotes a normal distribution with mean μ and standard deviation σ, and $\mathcal{U}(a,b)$ denotes a uniform distribution on the interval $[a,b]$;
2) weights $\boldsymbol{\alpha}$ (initialized with $\alpha_k = 1/(K^2-1) \forall k \neq \lfloor K^2/2 \rfloor$ in the first round);
3) $\sigma_{\mathcal{M}_1}$ (initialized with the signal's empirical standard deviation);
4) a normalizing assumption saying $\text{Prob}(y_i \in \mathcal{M}_1) = \text{Prob}(y_i \in \mathcal{M}_2)$.

In the M-step, vector $\boldsymbol{\alpha}$ is updated using a weighted least squares estimator

$$\boldsymbol{\alpha} = (\mathbf{Y}'\,\mathbf{W}\,\mathbf{Y})^{-1}\mathbf{Y}'\,\mathbf{W}\,\mathbf{y}. \tag{11}$$

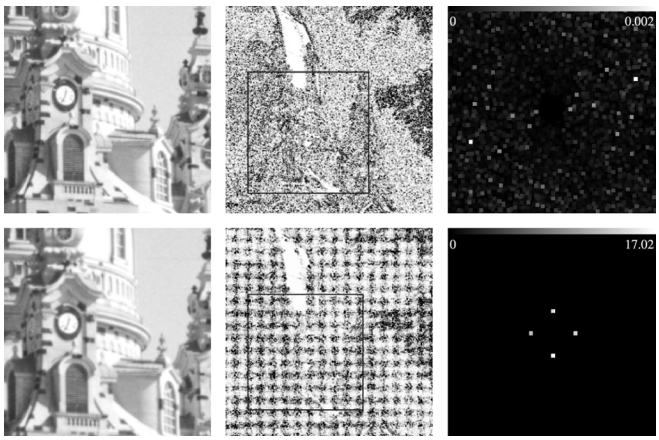

Fig. 2. Results of resampling detection for original image (top row) and 5% upsampling (bottom row). Complete p-maps are displayed in the middle column; frames mark the parts depicted on the left. Resampling artifacts lead to characteristic peaks in the corresponding spectrum (rightmost pictures).

Matrix $\mathbf{Y}$ has dimension $\dim(\mathbf{y}) \times K^2$ and contains the elements of all local neighborhoods as stacked row vectors (i.e., $\mathbf{Y} = ((\mathbf{N}^0 \cdot \mathbf{y})', (\mathbf{N}^1 \cdot \mathbf{y})', \ldots, (\mathbf{N}^{m_y n_y - 1} \cdot \mathbf{y})')$). Diagonal matrix $\mathbf{W}$ holds the corresponding conditional probabilities p_i of (10), hence $\mathbf{p} = \text{diag}(\mathbf{W})$. The new estimate for $\boldsymbol{\alpha}$ is used to update the local linear predictor $\mathbf{P}^{\boldsymbol{\alpha},i}$ and to calculate $\sigma_{\mathcal{M}_1}$ as a weighted standard deviation of the residuals ε

$$\sigma_{\mathcal{M}_1} = \sqrt{\frac{\sum_i p_i \cdot \varepsilon_i^2}{\sum_i p_i}}. \tag{12}$$

The E-step and M-step are iterated until convergence.

Resampling leaves a conspicuous periodical pattern in the so-obtained p-maps. This pattern becomes most evident in the frequency domain, using discrete Fourier transformation (DFT), where it shows up as distinct peaks that are typical for the resampling parameters. To enhance the visibility of the characteristic peaks, Popescu and Farid propose a contrast function γ [6]. It is composed of a radial weighting window, which attenuates low frequencies, and a gamma correction. The absolute values of the resulting complex plane can be visualized and presented to a human forensic investigator.

Fig. 2 illustrates the detection process. It compares an original grayscale image of size 350×350 to a processed version scaled up[1] with a linear interpolation to 105% of the original (left column). The resulting p-maps are displayed in the middle. As expected, the rather chaotic p-map of the original image turns to a salient periodic structure after transformation. This explains the different appearance of the spectrum (right column). To enhance the quality in print, each spectrum graph in this article is normalized to span the full intensity range. The range of spectral magnitudes is visualized by a gradient scale on top of each spectrum. We further apply a maximum filter to improve the visibility of the peaks.

[1]Upscaling is particularly prone to leave detectable traces in the redundancy of newly inserted pixels. So it forms a critical test for our methods.

V. Countermeasures Against Resampling Detection

In the hand of forensic investigators, Popescu and Farid's powerful detection method [6] might raise the temptation to use its results as proof of evidence in legal, social, and scientific contexts. However, one must bear in mind that forensic methods merely provide indications and are by orders of magnitude less dependable than other techniques, such as decent cryptographic authentication schemes. In contrast to cryptography, multimedia forensics remains an inexact science without rigourous security proofs. To draw attention to this problem, we will present different methods to perform image transformations that are almost undetectable by the aforementioned detector. We will identify some basic detection assumptions and describe methods to deliberately violate these assumptions. In this sense, the presented techniques can be considered as targeted attacks against the detection algorithm.

Prior to a detailed description of our attacks, we point out that we consider only never-compressed images. It is known that virtually all resampling detectors fail after moderate JPEG compression. The reason for this is twofold: First, periodic blocking artifacts interfere with periodic resampling artifacts.[2] Second, lossy compression blurs subtle periodic traces and, therefore, renders reliable detection impossible. Despite the existence of such a "universal" attack, we believe that research on targeted attacks against resampling detection is relevant. JPEG images are per se more likely to raise suspicion and, for example, news agencies may insist on never-compressed images.

A. Attacks Based On Nonlinear Filters

To detect suspicious periodic traces of previous resampling operations, the detector employs a local linear predictor, which expresses each pixel as the weighted sum of adjacent pixels (7). The detection method is therefore based on the assumption of linear dependences between pixels in a close neighborhood. Hence, all kinds of nonlinear filters $\xi : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, applied as the postprocessing step, are candidates for possible attacks (i.e., $\tilde{\mathbf{y}} = \xi(\mathbf{y})$). In this paper, we use the median filter $\xi = \xi_{\text{med}}$, a frequently used primitive operation in image processing [20], which replaces each pixel with the median of all pixels in a surrounding window of a defined shape and size. This acts as a lowpass filter, however, with floating cutoff frequency. Besides its nonlinear smoothing nature, median filtering seems an appropriate choice to serve as an attack against resampling detection as it is known to produce regions of constant or nearly constant intensity values [21]. Thus, median filtering destroys periodic dependencies between neighboring pixels, especially in homogeneous parts of the image.

Fig. 3 shows detection results for the 5% upscaled test image, postprocessed with a 5×5 median filter.[3] This straightforward attack can be called successful: periodic artifacts in the estimated p-map are largely suppressed, and consequently, the characteristic peaks in the spectrum have disappeared. However, at the same time, it becomes evident that the visual quality of the postprocessed image has suffered from noticeable blurring. This side effect can be attenuated using subtler (i.e., smaller) or more sophisticated filters (e.g., multistage median filters [22]), even though at the cost of higher detectability. Therefore, despite being effective in certain cases, the prospects of naive nonlinear filters for practical attacks remain limited.

Fig. 3. Results after upsampling by 5% and postprocessing with a 5×5 median filter: characteristic peaks in the spectrum vanish; however, the image appears excessively blurred.

B. Attacks Based on Geometric Distortion

A second basic assumption essential for successful resampling detection is the equidistance of the underlying sampling lattice. The detection method exploits the periodic structure in mapping the discrete lattice position from the source to the transformed image, in which a constant sequence of relative position of source and target pixels $\delta_{\mathbf{A}}(i,j)$ is repeated over the entire plane, cf (5). Periodic artifacts can be avoided if this systematic similarity is broken up.

Inspired by the effectiveness of geometric attacks against watermarking schemes [11], we have explored geometric distortion as the building block for attacks against resampling detection. To disrupt the specific similarity, each individual pixel's target position is computed from the transformation relation with a random disturbance vector $\mathbf{e} = (e_v, e_h)$ superimposed, that is

$$\tilde{R}_{i,j} = \phi\left(\delta_{\mathbf{A}}(i,j) + \mathbf{e}^j\right). \tag{13}$$

Horizontal and vertical distortion e_h and e_v are set to be noise samples, independently drawn from a zero mean Gaussian distribution $e_h, e_v \sim \mathcal{N}(0,\sigma)$. We use the notation $\mathbf{e}^j$ for the jth disturbance vector which displaces original pixel x_j and $\tilde{\mathbf{R}}$ for the modified resampling matrix (i.e., $\tilde{\mathbf{y}} = \tilde{\mathbf{R}} \cdot \mathbf{x}$). Parameter σ controls the strength of distortion.

Note that naive geometric distortion may leave strong visible jitter noise in the resulting image, which is perceived most visually disturbing at straight lines and edges. Horizontal distortion frays pronounced vertical structures and vice-versa. An extension of (13) evades such quality loss by modulating the strength of distortion adaptively from the local image content. More precisely, we apply two Sobel edge detectors [20] $\mathbf{S}^{\mathbf{h},i}$ for horizontal and $\mathbf{S}^{\mathbf{v},i}$ for vertical disturbance, respectively, to the transformed image $\mathbf{y}$ without any distortion

$$\tilde{R}_{i,j} = \phi\big(\delta_{\mathbf{A}}(i,j) + \mathbf{e}^j \odot (1 - \mathbf{S}^{\mathbf{h},i} \cdot \mathbf{y},\ 1 - \mathbf{S}^{\mathbf{v},i} \cdot \mathbf{y})\big). \tag{14}$$

Sobel filters are linear filters constructed equivalent to the predictor of (8), with a $3^2 \times \dim(\mathbf{y})$ neighborhood matrix and filter coefficients

$$\begin{aligned} \mathbf{h} &= (2^{-\ell} - 1) \otimes (1,0,-1,2,0,-2,1,0,-1) \text{ and} \\ \mathbf{v} &= (2^{-\ell} - 1) \otimes (1,2,1,0,0,0,-1,-2,-1) \end{aligned} \tag{15}$$

[2]This can be precluded to some degree (e.g., by suppressing characteristic JPEG frequencies in the p-map's spectrum).

[3]Magnified versions of the test image are depicted in Fig. 15.

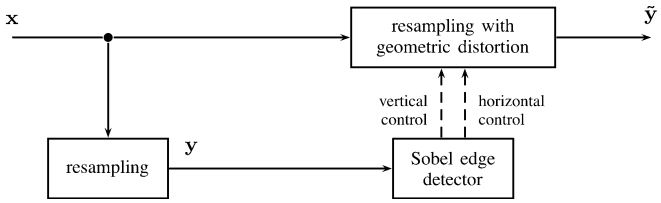

Fig. 4. Block diagram of geometric distortion with edge modulation.

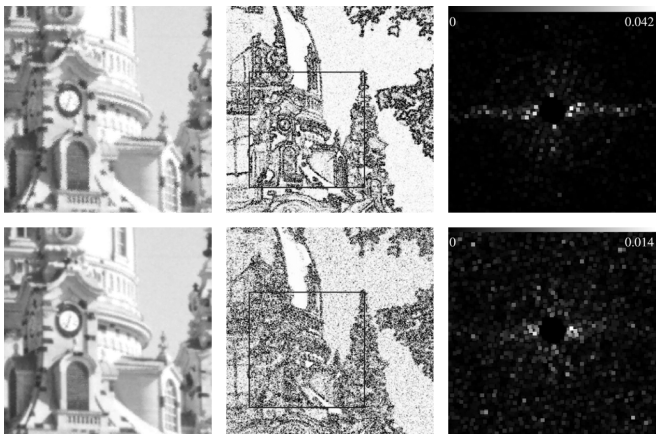

Fig. 5. Results after upsampling by 5% with geometric distortion of strength $\sigma = 0.4$. Comparison between naive distortion (top) and edge modulation using horizontal and vertical Sobel filters (bottom).

hence $\mathbf{S}^{*,i} \cdot \mathbf{y} \in [0, \ldots, 1]$. As pixels in the area of sharp edges will yield a large filter response, the modified attack construction ensures that less distortion is applied to regions where the visible impact would be most harmful otherwise. Since the filters are applied to an undistorted transformed image, this attack requires the image to be resampled twice, as depicted in the block diagram of Fig. 4.

Fig. 5 reports detection results for our 5% upscaled test image now using resampling with a geometric distortion of strength $\sigma = 0.4$. The results demonstrate that geometric distortion is capable of eliminating the characteristic traces from the p-map and its spectrum. In line with our expectations, edge modulation mitigates the loss in image quality considerably.

C. Dual-Path Approach to Undetectable Resampling

While geometric distortion with edge modulation generates already good results, we found from a comprehensive evaluation of many different transformation parameters that the undetectability can be further improved by applying different operations to the high- and low-frequency components of the image signal. Similar approaches have already been applied successfully for noise reduction [23] and attacks against spatial watermarks [24].

Adhering to a simple additive image model [20, Ch. 3], we use a dual-path approach to undetectable resampling. It combines median filtering (Section V-A) and geometric distortion (Section V-B) as depicted in the block diagram of Fig. 6. The image is modelled as sum of a low-frequency component and a high-frequency component $\mathbf{y} = \mathbf{y}^{(L)} + \mathbf{y}^{(H)}$. The two components are separated with a median filter $\mathbf{y}^{(H)} = \mathbf{y} - \xi_{\text{med}}(\mathbf{y})$.

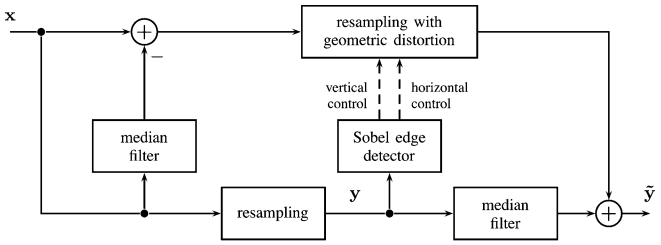

Fig. 6. Block diagram of the dual-path approach: combination of median filter for a low-frequency image component and geometric distortion with edge modulation for the high-frequency component.

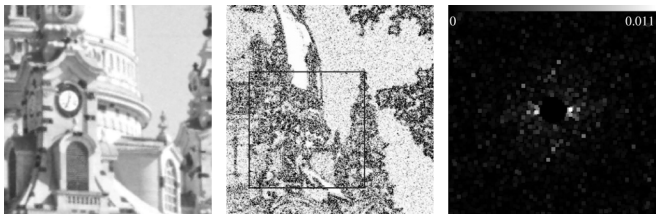

Fig. 7. Dual-path method: 5% upsampling, 7×7 median filter for a low-frequency component combined with geometric distortion ($\sigma = 0.3$) and edge modulation.

First, in the low-frequency path, the low-frequency component of the output image is obtained by applying a median filter directly to the resampled source image $\mathbf{y}$. Second, a high-frequency component is extracted from the source image $\mathbf{x}$ by subtracting the result of a median filter (other lowpass filters are conceivable as well). In the high-frequency path, this component is resampled with geometric distortion and edge modulation, where the edge information is obtained from the resampled image $\mathbf{y}$ prior to the median filter. The final image $\tilde{\mathbf{y}}$ is computed by adding up both components

$$\tilde{\mathbf{y}} = \xi_{\text{med}}(\mathbf{R} \cdot \mathbf{x}) + \tilde{\mathbf{R}} \cdot (\mathbf{x} - \xi_{\text{med}}(\mathbf{x})) \,. \tag{16}$$

This attack has two parameters: 1) the size of the median filter and 2) the standard deviation of the geometric distortion σ.

Fig. 7 reports the results of the dual-path approach for the 5% upscaled test image. Observe that the resulting p-map is most similar to the p-map of the original (see Fig. 2). Further, no suspicious peaks appear in its spectrum. The image quality is preserved and no artifacts are visible (cf. Fig. 15).

VI. Quantitative Evaluation

For a quantitative evaluation of our attacks against resampling detection, we use a database of 500 never-compressed 8-bit grayscale images. All images were taken with a Canon PowerShot S70 digital camera at full resolution (3112×2382 pixels) and stored in RAW format. In order to suppress possible interference from periodic patterns which might stem from color filter array (CFA) interpolation inside the camera [1], each image was downsampled by a factor two using nearest neighbor interpolation prior to any subsequent processing. This preprocessing was empirically found to be sufficient to reliably remove detectable traces of demosaicing by applying the CFA

detector [1] to the downsized images. We applied our attacks to a subset of 100 randomly chosen images, each resized and rotated by various degrees. More specifically, scaling matrices $\mathbf{A}^S$ with $0.5 \leq S \leq 2$ and S sampled in equidistant steps of $\Delta S = 0.05$ as well as rotation matrices $\mathbf{A}^\Theta$ with $0 < \Theta \leq \pi/4$ and Θ sampled in equidistant steps of $\Delta\Theta = \pi/36$ ($\hat{=} 5°$) were used for a total of 3900 resampled images per parameter setting. The remaining 400 images, further referred to as "training set," were used to determine the detector's decision threshold, cf. Section VI-B.

A. Performance Metrics

The performance metrics for the quantitative evaluation of the proposed attacks are twofold. First of all, the most relevant criterion is the (un)detectability of the conducted image transformations. We report detection rates (i.e., the fraction of correctly detected manipulations) for fixed false acceptance rates (FARs) of $\mathrm{FAR} \leq 1\%$ and $\mathrm{FAR} \leq 50\%$, respectively. Lower values indicate superior performance.

Any attempt to conceal resampling operations should not only be judged by the achieved level of undetectability, but also by the amount of image degradation in the resulting images compared to the resampled images with standard linear interpolation. For our quantitative evaluation, we chose two common image-quality metrics Q to assess the visual impact of the proposed attacks

$$Q = 20 \log \frac{2^\ell - 1}{\|(\mathbf{y} - \tilde{\mathbf{y}}) \odot \mathbf{v}\|}. \tag{17}$$

We report the metrics PSNR, where $\mathbf{v} = \mathbf{1}^{\dim(\mathbf{y})\times 1}$, and a variant adjusted for human visual perception wPSNR ("w" for weighted). It has been argued that the latter metric is a more valid indicator for the evaluation of attacks against watermarking schemes [12]. Weights $\mathbf{v}$ are computed from a so-called noise visibility function (NVF), which emphasizes image regions with high local variance and attenuates flat regions and soft gradients. Among the two NVFs proposed in [25], we have chosen the one based on a stationary generalized Gaussian image model [25, (26)]. Both metrics are measured in decibels. Higher values indicate superior image quality.

B. Automatic Detection of Resampling

As described in Section IV, the resampling detector relies on finding periodic dependencies between pixels in a close neighborhood. To identify forgeries automatically, Popescu and Farid propose measuring the similarity between the p-map of a given image and a set of synthetically generated periodic patterns [6]. They have found empirically that the synthetic map $\mathbf{s}^\mathbf{A}$ for transformation $\mathbf{A}$ can be obtained by computing the distance between each point in the resampled lattice and the closest point in the original lattice

$$s_i^\mathbf{A} = \left\| \mathbf{A} \cdot \nu_{m_s}^{-1}(i) - \left\lfloor \mathbf{A} \cdot \nu_{m_s}^{-1}(i) + \frac{1}{2} \otimes \mathbf{1}^{2\times 1} \right\rfloor \right\|. \tag{18}$$

As the detector lacks any prior knowledge about the actual transformation parameters $\mathbf{A}$, the detection process involves an exhaustive search in a sufficiently large set $\mathcal{A}$ of candidate transformation matrices.[4] In all of our experiments, we use a set of $|\mathcal{A}| = 692$ synthetic maps, 601 for scaling in the range of $0.5 \leq S \leq 2$, with S sampled in equidistant steps of $\Delta S = 0.0025$, and 91 for rotation in the range of $0 \leq \Theta \leq \pi/4$, with Θ sampled in equidistant steps of $\Delta\Theta = \pi/360$. The maximum pairwise similarity between an empirical p-map and all elements of $\mathcal{A}$ is taken as a decision criterion ρ [6]

$$\rho = \max_{\mathbf{A}\in\mathcal{A}} \left\| \left(\left| \gamma(\mathrm{DFT}(\mathbf{p})) \right| \odot \left| \mathrm{DFT}\left(\mathbf{s}^\mathbf{A}\right) \right| \right)^{1/2} \right\|. \tag{19}$$

Function γ is the contrast function (Section IV) and DFT is a 2-D DFT. If ρ exceeds a specific threshold ρ_T, then the corresponding image is flagged as resampled.[5] We have determined ρ_T empirically for defined false acceptance rates (cf. Section VI-A) by applying the detector to all 400 original images of the training set.

Note that the decision criterion ρ is not normalized with respect to the dimension of the analyzed image. In order to compare detection results from different resampling parameters, we always crop the center 256×256 block of the resampled image before it is presented to the detector.

C. Baseline Detection Results

To demonstrate the general effectiveness of the resampling detection scheme, Fig. 8 reports the detection rates for scaling (top) and rotation (bottom) using standard linear and bicubic interpolation (i.e., the case without attack). Each data point reflects 100 resampled images. The size of the detection neighborhood was set to $K = 5$. From the curves, we find perfect detection of upsampling and rotation, and still high accuracy for moderate downsampling. The decrease in detectability for small S is not surprising, as downsampling causes information loss whereas it is more difficult to impute new pixels with idiosyncratic information.[6] All in all, the results confirm a very reliable detection for a wide range of transformation parameters. Thus, Fig. 8 may serve as reference for the evaluation of our attacks with respect to their capability to hide traces of such image transformations.

D. Detectability of Sinc Interpolation

Before evaluating our targeted attacks against resampling detection, we study the detector's performance under sinc interpolation. Only recently, Mahdian and Saic have shown that, under an i.i.d. Gaussian signal model, ideal sinc interpolation will circumvent the formation of periodic artifacts in resampled images [16]. Interpolation with a sinc kernel, despite its higher computational demands, is therefore a critical benchmark for our attacks. Due to its infinite support, the sinc kernel is hard to

[4] To follow our notation of resampling (4)–(6), it is necessary to use synthetic scaling maps $\mathbf{s}^{\mathbf{A}^{1/S}}$ to obtain the corresponding periodic pattern for an actual transformation $\mathbf{A}^S$.

[5] This is a very conservative measure as cases may exist where $\rho > \rho_T$ although the best-matching synthetic map does not correspond to the actual transformation parameters. See Section VI-F for further discussion.

[6] As noted by Gallagher [14], phase-preserving upsampling by factor 2 (i.e., $\delta_\mathbf{A}(0,0) = (0.25, 0.25)$) is a sole exception that prevents the formation of periodic artifacts.

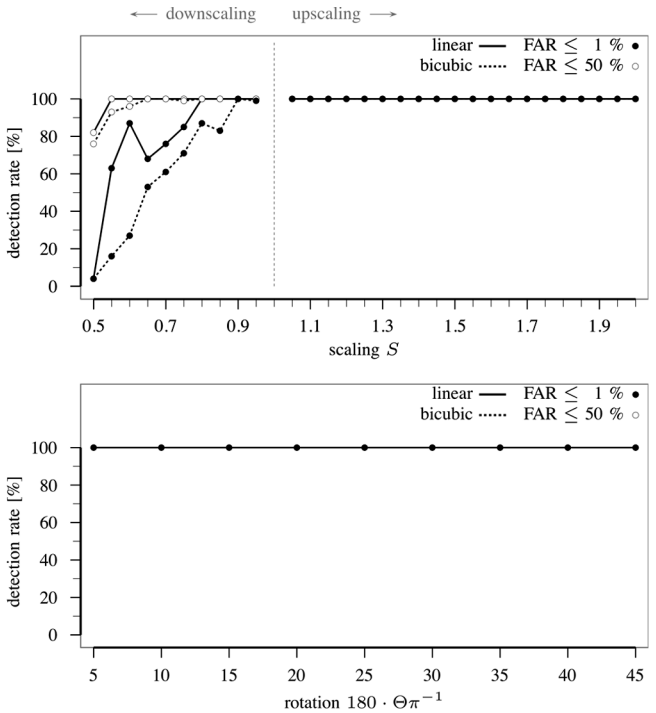

Fig. 8. Results of resampling detection after scaling (top) and rotation (bottom) by varying amounts for false acceptance rates $\mathrm{FAR} < 1\%$ and $\mathrm{FAR} < 50\%$, respectively. One-hundred resampled images are behind each datapoint.

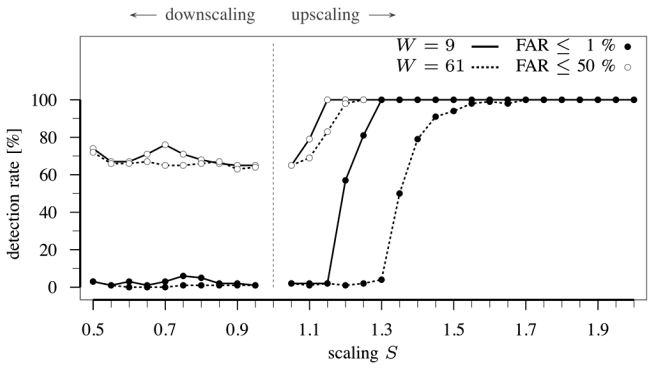

Fig. 9. Results of resampling detection after scaling using sinc-based interpolation. Downscaling and slight upscaling are virtually undetectable.

handle in practice. Typical implementations truncate it with an apodization function [26]. In our experiments, we used a rectangular window of finite support W. Fig. 9 reports detection rates for scaling and kernels with support $W = 9$ and $W = 61$. Observe that the detector indeed fails for downscaling and moderate upscaling, but stronger upscaling still remains perfectly detectable. The threshold which determines the transition from undetectable to highly detectable scaling depends on the support of the interpolation kernel. However, experiments have shown that even doubling the already large support of $W = 61$ yields no further gain in undetectability, which confirms the need for effective targeted attacks as an objective of this article. The observed deviation from the theoretical case might be due to the characteristics of typical image signals, which are, in general, not i.i.d. Gaussian.

Note that similar results can be achieved with higher order spline interpolation, which is asymptotically equivalent to sinc interpolation [27]. So we refrain from reporting more details.

E. Detectability of Median Filtering

Postprocessing with a median filter was introduced as an adequate attack in Section V-A. Fig. 10 reports detection rates (left) and average image quality (right) for scaled and median filtered images for filter sizes 3×3 and 7×7, respectively. While the detection rates remain on a relatively stable level of about 60% for the larger window size, the detectability for the 3×3 filtered images shows a strong dependence on the actual scaling factor S. Generally, larger window sizes introduce a higher degree of nonlinearity, resulting in less detectability in the upsampling case. Interestingly, the 3×3 filter is preferable for downsampled images, which follows from the comparatively "dense character" of downsized images. While larger windows tend to yield smoother images with a generally increased linear dependence between neighboring pixels, smaller windows will more likely choose pixels with idiosyncratic image content and, thus, lead to stronger local nonlinearities. However, the median filter has to be carefully chosen as larger windows cause substantial losses in image quality, which is observable in PSNR and wPSNR. Hiding traces of rotation typically requires smaller windows; a 5×5 filter already achieves detection rates below 25%.

We have also tested the detector's performance under multistage median filtering [22], which results in less blurred output images and increases image quality by up to 5 dB. However, only strong downscaling can be successfully concealed while upscaling remains detectable.

F. Detectability of Resampling With Geometric Distortion

Since for reasonable window sizes, median-filtered images may suffer from extensive blurring, we have investigated the effect of geometric distortion in the resampling process. Fig. 11 shows the results for scaling with distortion of strength $\sigma = 0.4$. They reveal a substantial gain in undetectability compared to the median filter. The curves indicate that edge modulation is beneficial in terms of visual quality and detectability. More specifically, detection rates with edge modulation are on average 6–20 p.p. below those without edge modulation. At the same time, the former yields an improvement in image quality (PSNR) of about 4 dB.

While the increase in visual quality intuitively follows from the signal adaptive modulation of the strength of distortion, the lowered detectability seems puzzling at first sight. To understand the observed decay, it is important to recall that the reported detection rates solely reflect the fraction of transformed images with $\rho > \rho_T$, independent of whether the correct transformation $\mathbf{A}$ (i.e., the synthetic map $\mathbf{s}^{\mathbf{A}}$) was selected or not. A synthetic map is considered as "correct" if the transformation parameters $\hat{\mathbf{A}}$ and true $\mathbf{A}$ do not differ by more than $10 \cdot \Delta S$ (i.e., tolerance ± 2.5 percentage points). A closer examination of the detection results suggests that for the majority of attacked images, the detector fails to find the correct transformation parameters. Table I exemplarily reports summary statistics for upsampling with $S = 1.4$. Observe that not a single "detection success" finds the correct synthetic map. Independent of the trans-

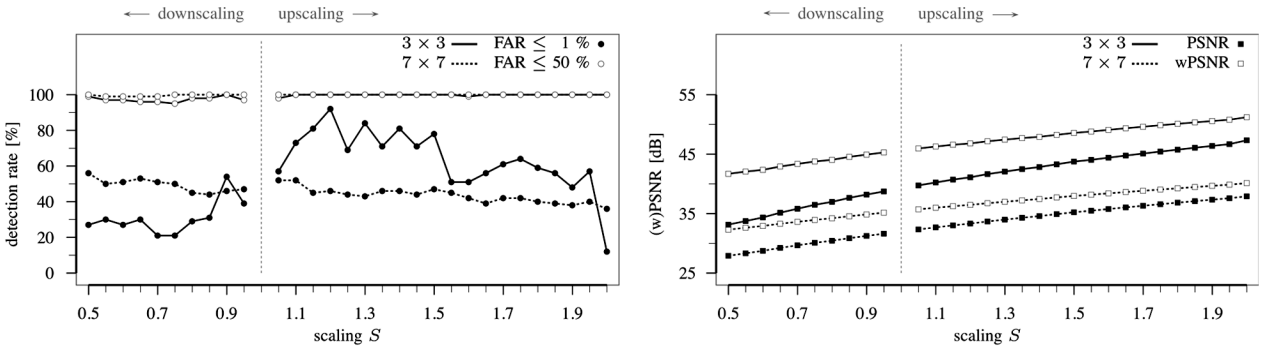

Fig. 10. Evaluation of a median filter at different window sizes. Detection rates (left) and average image quality (right). Larger window sizes reduce detection rates in the upscaling case; however, a small window is preferable for downscaling. Small windows retain higher image quality.

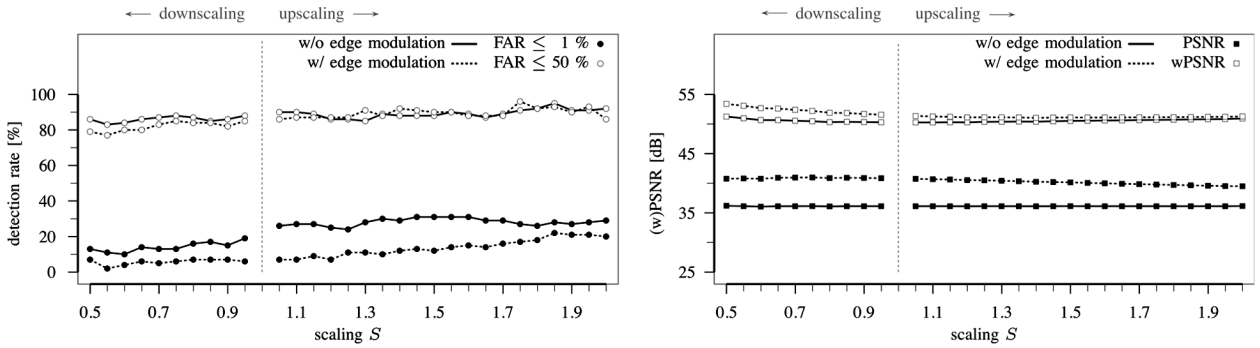

Fig. 11. Evaluation of geometric distortion ($\sigma = 0.4$) with and without edge modulation. Detection rates (left) and image quality (right). Edge modulation yields substantially better quality and superior detection results.

TABLE I
BREAKDOWN OF DETECTION DECISION

	resampling with geometric distortion ($S = 1.4$)			
	w/o Sobel		w/ Sobel	
	$\rho \leq \rho_T$	$\rho > \rho_T$	$\rho \leq \rho_T$	$\rho > \rho_T$
false map	73	26	70	12
corr. map	1	0	18	0
Σ	74	26	88	12

formation parameters, p-maps smoothed by our attacks happen to fit to some candidate synthetic maps with high amplitude in the low-frequency coefficients. Plain geometric distortion creates smoother p-maps due to a larger stochastic support and, therefore, is more susceptible to such "false map" alarms.

G. Detectability of the Dual-Path Approach

Finally, Fig. 12 presents the results for the dual-path approach for scaling (top row) and rotation (bottom row). As for resampling with geometric distortion (Section VI-F), the strength of distortion was set to $\sigma = 0.4$. The frequency components have been separated with 5×5 and 7×7 median filters, respectively. For the chosen strength of distortion, the dual path generally yields very low detection rates of less than 20% for all tested resampling parameters ($\mathrm{FAR} \leq 1\%$).

The curves indicate that the filter size is not crucial with respect to detectability (left column). However, smaller windows might be preferred when image quality matters (right column). From a comprehensive evaluation of different attack settings, we have found that the level of undetectability is largely determined by the strength of geometric distortion σ. Fig. 14 reports detection rates ($\mathrm{FAR} \leq 1\%$) for sample scaling factors at varying strengths of geometric distortion. The size of the median filter was fixed to 5×5. Roughly speaking, the detector's sensitivity to geometric distortion can be well approximated by a step function. While too little distortion has practically no influence on detectability, too much distortion is not rewarded with better undetectability. The minimum distortion necessary to yield satisfactory results slightly varies for different resampling parameters. However, $\sigma = 0.4$ seems to be a reasonably safe default. It is important to note that a careful choice of σ is essential to maintain acceptable image quality, because stronger distortion will cause visually more perceivable image degradation.

H. Finding the Conditional Best Hiding Method

A direct comparison of the dual-path approach with geometric distortion as described in Section V-B reveals an advantage of the former especially for upscaling. Fig. 13 depicts comparative results for $\sigma = 0.4$ in terms of the achieved gain in undetectability (left, solid line) and image quality metrics (right). The curves indicate that, on average, the advantage of the dual-path approach increases with the scaling factors ($S \gg 1$) up to ten percentage points at the cost of only marginal compromises in image quality ($< 3\,\mathrm{dB}$). Nevertheless, for downscaling (and rotation, not printed due to space constraints), plain geometric distortion appears to perform

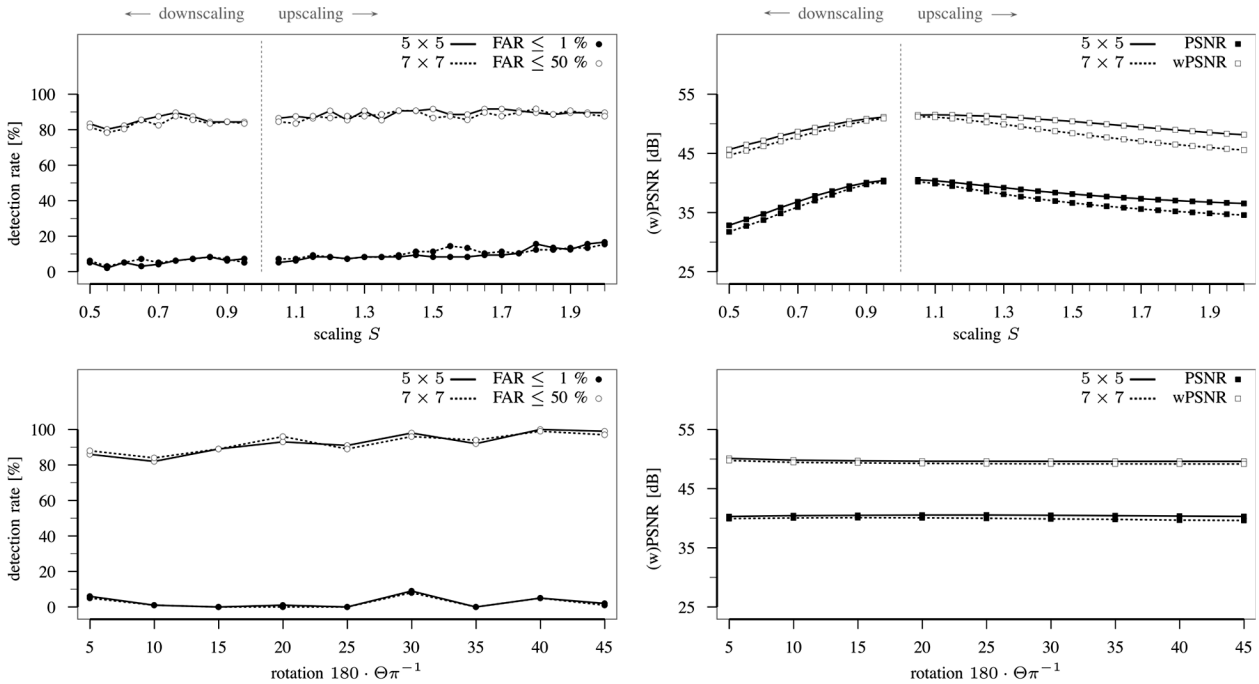

Fig. 12. Evaluation of dual-path approach for scaling (top row) and roation (bottom row). Detection rates (left column) and average image quality (right column) for $\sigma = 0.4$. The breakdown by window size of the median filter (5×5 versus 7×7) and false acceptance rates (FAR: 1% versus 50%). Observe the very low detection rates independent of the window size for all resampling parameters. Smaller windows sizes in the low-frequency component retain better image quality.

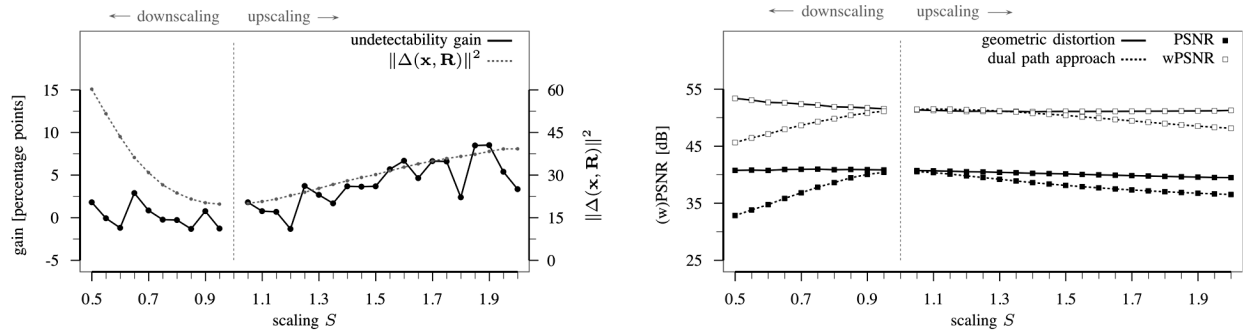

Fig. 13. Comparison of geometric distortion and dual-path approach for scaling. Undetectability gain (left) and image quality (right). $\sigma = 0.4$; 5×5 median filter; $\mathrm{FAR} \leq 1\%$.

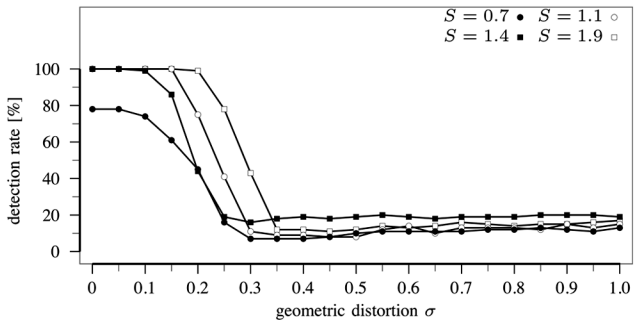

Fig. 14. Dual-path approach: detectability versus strength of distortion for sample scaling factors; 5×5 median filter; $\mathrm{FAR} \leq 1\%$.

better with comparable detection rates, but notably better image quality.

To further study the detectability differential between geometric distortion and dual-path approach, we rewrite (16) as

$$\hat{\mathbf{y}} = \tilde{\mathbf{R}}\mathbf{x} + \xi_{\mathrm{med}}(\mathbf{R}\mathbf{x}) - \tilde{\mathbf{R}}\xi_{\mathrm{med}}(\mathbf{x}) = \tilde{\mathbf{R}}\mathbf{x} + \Delta(\mathbf{R},\mathbf{x}) \quad (20)$$

where $\tilde{\mathbf{R}}$ is the modified resampling matrix of (14) (i.e., function $\Delta(\mathbf{R},\mathbf{x})$ returns the difference signal between an image resampled with geometric distortion and the dual-path approach). Larger relative scaling factors $|1-S|$ increase the difference between signals $\xi_{\mathrm{med}}(\mathbf{R}\mathbf{x})$ and $\tilde{\mathbf{R}} \cdot \xi_{\mathrm{med}}(\mathbf{x})$. As a result, the energy of the noise signal $\Delta(\mathbf{R},\mathbf{x})$ increases with $|1-S|$. Fig. 13 (left) includes a curve of average energy measured by the square sum of differences $\|\Delta(\mathbf{R},\mathbf{x})\|^2$, which grows with the increasing performance advantage of the dual-path approach for $S > 1$. For downscaling, however, $\|\Delta(\mathbf{R},\mathbf{x})\|^2$ grows even faster with increasing $1-S$ although the performance of the dual-path approach does not improve compared to resampling with geometric distorton (but image quality deteriorates due to the specific noise of the dual-path approach). Therefore, we conjecture that the dual-path approach behaves very similar to resampling with geometric distortion for moderate scaling and benefits from its specific signal- and transformation-adaptive "postprocessing" only for larger relative scaling factors.

Fig. 15. Visual comparison: test image and image-quality metrics after resampling to 105%. From left to right: plain linear interpolation, 5×5 median filtering, geometric distortion with edge modulation ($\sigma = 0.4$), dual-path approach (7×7 median $\sigma = 0.3$).

I. Robustness and Validation

Note that we have also tested the robustness of our results for detectors with smaller $(K = 3)$ and larger $(K = 7)$ neighborhoods. As the corresponding detection rates do not differ substantially from the reported figures, we conclude that our results are fairly robust with regard to Popescu and Farid's detector and refrain from reporting them separately.

To validate the effectiveness of the proposed attacks against related resampling detectors, we implemented Mahdian and Saic's derivative-based method [16], the most relevant alternative. It turned out that it is very difficult to find appropriate decision thresholds for our images. For plain upsampling, we found false acceptance rates as high as 30% at 100% correct detection. These results do not match the performance of Popescu and Farid's method. As aggregated graphs for $\mathrm{FAR} < 1\%$ are not very indicative, we compared the continuous score for individual images resized with and without our attacks. In almost every case, the score was lower after the attack. We therefore conjecture that derivative-based resampling detection will also fail on images resampled with our attacks.

VII. Conclusion

This paper has taken a critical view on the reliability of forensic techniques as tools to generate evidence of authenticity for digital images. In particular, we have pointed out how tamper-hiding techniques, in general, can be integrated in a broader ontology of multimedia security disciplines. The main contribution of this paper is a presentation and evaluation of three approaches to defeat a specific method of resampling detection, which has been developed to unveil scaling and rotation operations of digital images or parts thereof. These attacks have turned out to be the most effective ones in a broader research effort, which also led to a number of dead ends. Some of the alternative attack methods are briefly documented in [9]. Among the successful methods, resampling with edge-modulated geometric distortion (for downsampling and rotation) and the dual-path approach (primarily for upsampling), which complements the former by a median filter of the low-frequency component of the image signal, achieved the best performance and should be regarded as benchmarks for other specific tamper-hiding techniques. At the same time, we would like to point out that the resampling detector of Popescu and Farid [6], against which our work in this article is targeted, is, to the best of our knowledge, the most reliable detector of standard interpolation. We have selected this particular detector with the aim of building a sample attack against a powerful and challenging method.

Apart from the detailed results presented so far, there are at least two more general conclusions worth mentioning. First, attacks which are integrated in the manipulation operation appear to be more effective than others that work at a post-processing step. This is plausible since information about the transformation parameters is not available at the postprocessing stage. Therefore, much stronger interference with the image structure is necessary to cover up statistical traces of all possible transformations. Second, a closer look at all quantitative results suggests that downscaling and rotation are easier to conceal than upscaling. This is plausible too, since downscaling causes information loss, whereas it is more difficult to impute new pixels with idiosyncratic information. This implies that larger window sizes (for the median filter) and stronger geometric distortion are necessary for upscaling to achieve similar levels of (un)detectability as for downscaling.

As for the limitations, we consider this paper to be an early and modest attempt in an interesting subfield. It is obvious that our results hold only for the specific class of detectors and we cannot rule out that image manipulations conducted with our proposed methods are detectable with 1) other existing forensic techniques or 2) new targeted detection methods that are build with the intention to discover our attacks. While this might trigger a new cat-and-mouse race between forensic and counter-forensic techniques, we believe that such creative competition is fruitful and contributes to a more holistic picture on the possibilities and limitations of image forensics, an area where much prior research has been done against the backdrop of a fairly naive "adversary model"—a term borrowed from cryptography, where anticipating strong and knowledgeable adversaries has a longer tradition [28].

More research questions are abundant: It would be desirable to have a formal framework that explains why the building blocks (median filter, geometric distortion) effectively suppress periodic linear residuals and how they interact. This could lead to a theory that allows deriving the best method conditional to a larger parameter space than explorable with experiments. In addition, each building block could be optimized separately (e.g.,

replacing the median filter for the dual-path decomposition by a more tailored filter, or introducing a weighted synthesis of the components to minimize distortion). Further issues emerge if local resampling in parts of larger images shall be concealed [29], or when generalizing the methods to color images where plausible CFA patterns have to be introduced.

On a more abstract level, one may ask the question as to whether it is possible at all to construct provable secure techniques under gentle assumptions. We conjecture that an ultimate response is far distant and is probably linked to related "hard problems," such as the search for provable secure high-capacity steganography (with realistic cover assumptions), and to the development of much better stochastic image models.

ACKNOWLEDGMENT

The authors would like to thank their colleagues S. Berthold, T. Gloe, and S. Köpsell for their comments. The database of test images was provided by T. Gloe.

REFERENCES

[1] A. C. Popescu and H. Farid, "Exposing digital forgeries in color filter array interpolated images," *IEEE Trans. Signal Process.*, vol. 53, no. 10, pp. 3948–3959, Oct. 2005.

[2] M. K. Johnson and H. Farid, "Exposing digital forgeries through chromatic aberration," in *Proc. Multimedia and Security Workshop*, 2006, pp. 48–55.

[3] M. Chen, J. Fridrich, J. Lukáš, and M. Goljan, "Imaging sensor noise as digital X-ray for revealing forgeries," in *Information Hiding*, ser. Lect. Notes Comput. Sci. 4567, T. Furon, F. Cayre, G. Doërr, and P. Bas, Eds. Berlin, Germany: Springer Verlag, 2007, pp. 342–358.

[4] A. Swaminathan, M. Wu, and K. J. R. Liu, "Digital image forensics via intrinsic fingerprints," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 101–117, Mar. 2008.

[5] J. Fridrich, D. Soukal, and J. Lukáš, "Detection of copy-move forgery in digital images," in *Digital Forensic Research Workshop*, 2003.

[6] A. C. Popescu and H. Farid, "Exposing digital forgeries by detecting traces of re-sampling," *IEEE Trans. Signal Process.*, vol. 53, no. 2, pp. 758–767, Feb. 2005.

[7] H. Farid, "Exposing digital forgeries in scientific images," in *Proc. Multimedia and Security Workshop*, 2006, pp. 29–36.

[8] M. K. Johnson and H. Farid, "Exposing digital forgeries in complex lighting environments," *IEEE Trans. Inf. Forensics Security*, vol. 2, no. 3, pp. 450–461, Sep. 2007.

[9] M. Kirchner and R. Böhme, "Tamper hiding: Defeating image forensics," in *Information Hiding*, ser. Lect. Notes Comput. Sci. 4567, T. Furon, F. Cayre, G. Doërr, and P. Bas, Eds. Berlin, Germany: Springer-Verlag, 2007, pp. 326–341.

[10] R. J. Anderson and F. A. P. Petitcolas, "On the limits of steganography," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 4, pp. 474–481, May 1998.

[11] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Attacks on copyright marking systems," in *Information Hiding*, ser. Lect. Notes Comput. Sci. 1525, D. Aucsmith, Ed. Berlin, Germany: Springer-Verlag, 1998, pp. 219–239.

[12] S. Voloshynovskiy, S. Pereira, A. Herrigel, N. Baumgaertner, and T. Pun, "Generalized watermarking attack based on watermark estimation and perceptual remodulation," in *Proc. SPIE: Security and Watermarking of Multimedia Content II*, P. W. Wong and E. J. Delp, Eds. San Jose, CA: SPIE, 2000, vol. 3971, pp. 358–370.

[13] A. Piva and M. Barni, "The first BOWS contest (break our watermarking system)," in *Proc. SPIE: Security and Watermarking of Multimedia Content IX*, E. J. Delp and P. W. Wong, Eds. San Jose, CA: SPIE, 2007, vol. 6505, 650516.

[14] A. C. Gallagher, "Detection of linear and cubic interpolation in JPEG compressed images," in *Proc. 2nd Canadian Conf. Computer and Robot Vision*, 2005, pp. 65–72.

[15] S. Prasad and K. Ramakrishnan, "On resampling detection and its application to detect image tampering," in *Proc. Int. Conf. Multimedia Expo.*, 2006, pp. 1325–1328.

[16] B. Mahdian and S. Saic, "Blind authentication using periodic properties of interpolation," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 3, pp. 529–538, Jul. 2008.

[17] A. C. Popescu, "Statistical tools for digital image forensics," Ph.D. dissertation, Dept. Comput. Sci., Dartmouth College, Hanover, NH, 2005.

[18] G. Wolberg, *Digital Image Warping*, 3rd ed. Los Alamitos, CA: IEEE Comput. Soc. Press, 1994.

[19] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *J. Roy. Stat. Soc., Ser. B*, vol. 39, no. 1, pp. 1–38, 1977.

[20] I. Pitas, *Digital Image Processing Algorithms and Applications*. New York: Wiley, 2000.

[21] A. C. Bovik, "Streaking in median filtered images," *IEEE Trans. Acoust., Speech Signal Process.*, vol. ASSP-35, no. 4, pp. 493–503, Apr. 1987.

[22] G. R. Arce and R. E. Foster, "Detail-preserving ranked-order based filters for image processing," *IEEE Trans. Acoust., Speech Signal Process.*, vol. 37, no. 1, pp. 83–98, Jan. 1989.

[23] R. Bernstein, "Adaptive nonlinear filters for simultaneous removal of different kinds of noise in images," *IEEE Trans. Circuits Syst.*, vol. CAS-34, no. 11, pp. 1275–1291, Nov. 1987.

[24] G. C. Langelaar, J. Biemond, and R. L. Lagendijk, "Removing spatial spread spectrum watermarks by non-linear filtering," *Proc. EUSIPCO*, pp. 2281–2284, 1998.

[25] S. Voloshynovskiy, A. Herrigel, N. Baumgaertner, and T. Pun, "A stochastic approach to content adaptive digital image watermarking," in *Information Hiding*, ser. Lect. Notes Comput. Sci. 1768, A. Pfitzmann, Ed. Berlin, Germany: Springer-Verlag, 2000, pp. 212–236.

[26] P. Thévenaz, T. Blu, and M. Unser, "Interpolation revisited," *IEEE Trans. Med. Imag.*, vol. 19, no. 7, pp. 739–758, Jul. 2000.

[27] A. Aldroubi, M. Unser, and M. Eden, "Cardinal spline filters: Stability and convergence to the ideal sinc interpolator," *Signal Process.*, vol. 28, no. 2, pp. 127–138, 1992.

[28] A. Kerckhoffs, "La cryptographie militaire," *J. Sci. Mil.*, vol. IX, pp. 5–38, 1883.

[29] M. Kirchner, "On the detectability of local resampling in digital images," in *Proc. SPIE: Security and Watermarking of Multimedia Content IX*, E. J. Delp and P. W. Wong, Eds. San Jose, CA: SPIE, 2008, vol. 6819, 68190.

Matthias Kirchner graduated in information systems technology from Technische Universität Dresden, Dresden, Germany, in 2007, where he is currently pursuing the Ph.D. degree in the privacy and security group. He received a doctorate scholarship from Deutsche Telekom Stiftung, Bonn, Germany.

His particular interests include multimedia forensics and information hiding.

Rainer Böhme received the M.A. degree in communication science, economics, and computer science from Technische Universität Dresden, Dresden, Germany, where he is currently pursuing the Ph.D. degree in the privacy and security group.

His research interests include steganography, steganalysis, and multimedia forensics as well as the economics of privacy and information security.