Linear Row and Column Predictors for the Analysis of Resized Images

Matthias Kirchner
Institute of Systems Architecture
Technische Universität Dresden
01062 Dresden, Germany
matthias.kirchner@inf.tu-dresden.de

ABSTRACT

This paper adds a new perspective to the analysis and detection of periodic interpolation artifacts in resized digital images. Instead of relying on a single, global predictor, we discuss how the specific structure of resized images can be explicitly modeled by a series of linear predictors. Characteristic periodic correlations between neighboring pixels are then measured in the estimated predictor coefficients itself. Experimental results on a large database of images suggest a superior detection performance compared to state-of-the-art methods.

Categories and Subject Descriptors

I.4.m [**Image Processing**]: Miscellaneous

General Terms

Algorithms, Security

Keywords

digital image forensics, tamper detection, resampling detection, linear prediction, periodic artifacts

1. INTRODUCTION

Resampling detection [12, 4, 10, 7], meanwhile a standard problem in digital image forensics [13, 3], strives to analyze and expose traces of geometric image transformations such as scaling or rotation. These image processing primitives are of particular interest in the forensic analysis of digital images for several reasons. First and foremost, complex image forgeries oftentimes rely on the transformation of (parts of) images when objects have to be aligned in their size or perspective. Moreover, while especially the resizing of digital images can be seen as a type of plausible post-processing, which is in the first instance without consequences on the image's integrity, it is still of relevance in a general forensic analysis of the image, where every information about the processing history adds value. In this context, a detection of downscaling is particularly important, as a reduction in image size is a common form of universal attack to hide traces of previous manipulations [2].

Most of the prior work in the literature that is concerned with a detection of resampling exploits periodic linear correlations between neighboring pixels. These correlations are an inevitable side-product of interpolation, which itself is the key to visually appealing image transformations. In this paper, we follow this general approach. However, we add a new perspective and thereby hope to deepen the understanding of the specific artifacts in resized images. In particular, we will show how a series of tailored linear predictors, which anticipate the inherent structure of resized images, can be used to explicitly model and detect traces of interpolation. This goes beyond prior work, where interpolation artifacts are usually measured in the residue of one global predictor (or equivalently in a difference signal obtained by linear filtering with a fixed kernel).

The remainder of this paper starts with an overview of interpolation artifacts in resized images in Sect. 2, before Sect. 3 presents our findings on how the characteristics of resized images can be exploited to construct a suitable detector. Section 4 then reports experimental results from a large database of images and motivates some concluding remarks in Sect. 5.

2. INTERPOLATION ARTIFACTS

In typical signal processing applications, interpolation is implemented by means of a linear kernel that assigns scalar weights to genuine samples in close proximity. Since virtually all n-dimensional interpolation kernels of practical relevance are of separable nature, we will restrict ourselves in the following to one-dimensional (1D) signals whenever applicable. Then, interpolation of a signal s at an arbitrary real-valued position $x \in \mathbb{R}$ with a kernel $h\colon \mathbb{R} \to \mathbb{R}$ is to be written as a linear combination of original samples at integer positions $\chi \in \mathbb{Z}$,

$$s(x) = \sum_{\chi=-\infty}^{\infty} h(x-\chi)s(\chi)\,. \tag{1}$$

Typically, the need for interpolation is driven by a specific geometric transformation, which determines the mapping of integer source coordinates χ' to real-valued target coordinates x. Resizing, or scaling—the most fundamental mapping—links target coordinates with source coordinates through a linear factor ω, $x = \omega\chi'$. The factor ω is usually

MM&Sec'10, September 9–10, 2010, Roma, Italy.

referred to as the inverse scaling factor, where $\omega < 1$ corresponds to upsampling and $\omega > 1$ means downsampling, respectively.

To recall that a scaling factor $\omega^{-1} = p/q$, with p and q relatively prime, $p \perp q$, results in periodic interpolation artifacts with period p, we firstly note that interpolation weights $h(x-\chi)$ can, without loss of generality, be expressed in terms of the differences $\delta_x = x - \lfloor x \rfloor$ and rewrite (1) as

$$\sum_{\chi=-\infty}^{\infty} h(x-\chi)s(\chi) = \sum_{l=-\infty}^{\infty} h\left(l-\delta_x\right) s(\lfloor x \rfloor + l)\,. \qquad (2)$$

From (2) follows that two samples $s(x_1)$ and $s(x_2)$ are interpolated from their corresponding neighbors with exactly the same set of interpolation coefficients whenever $\delta_{x_1} = \delta_{x_2}$. By further noting that scaled discrete coordinates $x = \omega\chi'$, with $\omega = q/p$ and $p \perp q$, can be written as

$$x = \frac{q}{p}\left(\frac{pm}{q} + n\right), \quad \text{with } m \in \mathbb{Z} \text{ and } n \in \{0, 1, \ldots, p-1\},$$

the periodicity of δ_x with period p becomes apparent:

$$\delta_x = m + \frac{qn}{p} - \left\lfloor m + \frac{qn}{p} \right\rfloor = \frac{qn}{p} - \left\lfloor \frac{qn}{p} \right\rfloor .$$

The periodicity of interpolation weights ultimately leads to periodic linear correlations between neighboring interpolated samples. The specific form of correlation generally depends on the difference δ_x. Assuming a symmetric auto-correlation in the original signal as well as a symmetric interpolation kernel, a value $0 < \delta_x < 0.5$ means that the corresponding sample is stronger correlated with its left neighbors (and vice versa for $0.5 < \delta_x < 1$). Both $\delta_x = 0$ and $\delta_x = 0.5$ result in a symmetric dependence on the right and left neighbors, respectively. Since δ_x is monotonically increasing within one period, scaled signals typically exhibit an asymmetric auto-correlation that has a gradually varying, periodic bias.

It was Popescu and Farid's seminal paper [11, 12] that first discussed the existence of such characteristic correlations as an indication of geometric transformations. The authors proposed a detector based on a linear predictor. By predicting each sample from its neighbors, the residual signal e,

$$e(x) = e(\omega\chi) = s(\omega\chi) - \sum_{\substack{|k|\le K \\ k\neq 0}} \alpha_k \; s(\omega\chi + \omega k)\,, \qquad (3)$$

gives information about strength and characteristics of the linear dependencies between neighboring samples. Large absolute differences indicate a minor degree of linear correlation and vice versa. In this setting, interpolation generally leads to periodic artifacts in the predictor residue [7].

In preparation for our following considerations, we would like to stress that the very same predictor coefficients $\boldsymbol{\alpha}$ are applied for all samples throughout the entire signal. Equation (3) can thus be understood as a convolution of the interpolated signal with a fixed linear kernel with predictive properties. Popescu and Farid devised a weighted least squares (WLS) estimation of the kernel coefficients, which they embedded in an expectation/maximization framework. The estimation procedure chooses those coefficients that minimize the overall weighted squared prediction error, with the weights being determined in an iterative process [12]. Ultimately, the estimated coefficients can be interpreted as a measure of *average* linear dependence between neighboring samples. It is thus not surprising to see that subsequent work reported a strong symmetry ($\alpha_k \approx \alpha_{-k}$) and, to a certain degree, invariance with respect to the parametrization of the underlying geometric transformation [7].

A largely equivalent approach for the detection of interpolation artifacts stems from the analysis of (partial) derivatives. As shown by Gallagher [4], and later on generalized by Mahdian and Saic [10], the variance of the (n-th order) derivatives of interpolated signals exhibits a characteristic periodic pattern. Interesting parallels to predictor-based detectors become apparent by recalling that typical implementations of discrete derivative operators are based on a convolution with a fixed symmetric kernel [7].

In summary, we can conclude that, to the best of our knowledge, the vast majority of resampling detectors uncovers interpolation artifacts in some type of residue signal, which is obtained by linear filtering with a fixed kernel. In the following, we will vary this procedure by not only employing a single (i. e., global) predictor, but rather a series of local predictors that anticipate the characteristic structure of resized images. As a consequence, periodic interpolation artifacts can be measured and analyzed more explicitly, as they manifest themselves in the actual predictor coefficients.

3. PREDICTOR COEFFICIENTS TO MEASURE INTERPOLATION ARTIFACTS

The key observation of this paper follows from the separability of typical interpolation kernels. For resized images, a separable interpolation kernel means that all pixels within one row are correlated with their corresponding vertical neighbors in exactly the same way (the same holds for columns and their horizontal neighbors).[1] This characteristic structure can be expressed in terms of a linear model by denoting $\boldsymbol{r}^{(i)}$ as the vector of all pixels in the i-th row,

$$\boldsymbol{r}^{(i)} - \boldsymbol{\eta}^{(i)} = \sum_{\substack{|k|\le K \\ k\neq 0}} \alpha_k^{(i)} \left(\boldsymbol{r}^{(i-k)} - \boldsymbol{\eta}^{(i-k)}\right) + \boldsymbol{\epsilon}^{(i)}\,. \qquad (4)$$

In the above model, $\boldsymbol{\eta}$ means quantization errors due to rounding and $\boldsymbol{\epsilon}$ describes the model error, i. e. the portion of $\boldsymbol{r}^{(i)} - \boldsymbol{\eta}^{(i)}$ that cannot be explained by linear inter-pixel dependencies. By assuming a suitable neighborhood size K, the model error vanishes for upsampling ($\omega < 1$), i. e., each pixel is (apart from rounding errors) fully determined by its vertical neighbors. For downsampling, this is not necessarily the case and rather depends on the support of the interpolation kernel and the inverse scaling factor ω.

The row index in (4) highlights the difference in our approach of modeling interpolation artifacts compared to prior art. Each row has its own predictor $\boldsymbol{\alpha}^{(i)}$ that explicitly reflects the actual linear correlation between pixels in that row and their vertical neighbors. Contrary to an implicit model of periodic artifacts by means of the deviation from the average predictor, a series of tailored predictors allows to measure periodic traces in the predictor coefficients itself. The estimation of the predictor coefficients can be conducted using standard linear regression methods, whereas one re-

[1] For notational convenience, we will subsequently only refer to rows and their vertical neighbors. The reader should however keep in mind that our considerations can be readily reframed to a column-wise model.

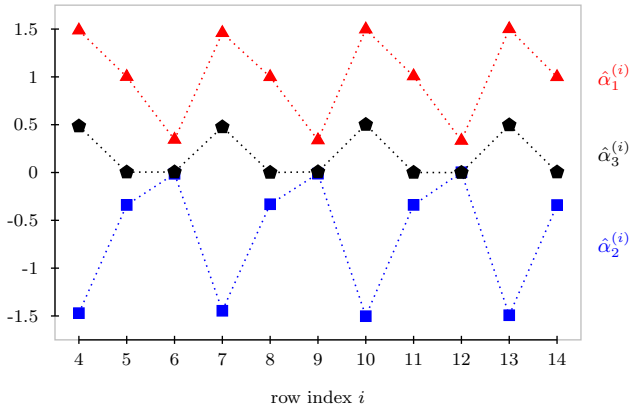

Figure 1: TLS estimates of row predictor coefficients for a sequence of 11 rows of a 150 % bilinearly upsampled image ($K = 3$, image size 350×350 pixel). The expected coefficients are $\{3/2, 1, 1/3\}$ for α_1, $\{-3/2, -1/3, 0\}$ for α_2 and $\{1/2, 0, 0\}$ for α_3.

gression is required per row. The following subsections shall discuss concrete implementations.

3.1 Ideal case: total least squares

For a vanishing model error, $\boldsymbol{\epsilon}^{(i)} = \mathbf{0}$, all pixels of a particular row can (apart from rounding errors) be fully described as a linear combination of their vertical neighbors. Since in (4), both dependent and explanatory variables are subject to measurement errors, a total least squares (TLS) estimation of the predictor coefficients is appropriate [14].

Rewriting (4) as

$$\boldsymbol{r}^{(i)} - \boldsymbol{\eta}^{(i)} = \left(\mathbf{R}^{(i)} - \mathbf{N}^{(i)}\right) \cdot \boldsymbol{\alpha}^{(i)} ,$$

$$\text{with} \quad \mathbf{R}^{(i)} = \left[\boldsymbol{r}^{(i-K)}, \ldots, \boldsymbol{r}^{(i-1)}, \boldsymbol{r}^{(i+1)}, \ldots, \boldsymbol{r}^{(i+K)}\right]$$

$$\text{and} \quad \mathbf{N}^{(i)} = \left[\boldsymbol{\eta}^{(i-K)}, \ldots, \boldsymbol{\eta}^{(i-1)}, \boldsymbol{\eta}^{(i+1)}, \ldots, \boldsymbol{\eta}^{(i+K)}\right] ,$$

TLS aims at finding a coefficient vector $\hat{\boldsymbol{\alpha}}^{(i)}$, such that

$$\hat{\boldsymbol{r}}^{(i)} = \hat{\mathbf{R}}^{(i)} \cdot \hat{\boldsymbol{\alpha}}^{(i)} \quad \text{and} \quad \left\| \left[\mathbf{R}^{(i)}, \boldsymbol{r}^{(i)}\right] - \left[\hat{\mathbf{R}}^{(i)}, \hat{\boldsymbol{r}}^{(i)}\right] \right\|_F \to \min .$$

The solution follows from the singular value decomposition (SVD) of the augmented matrix $\left[\mathbf{R}^{(i)}, \boldsymbol{r}^{(i)}\right]$, cf. [14].

In general, TLS allows a very precise estimation of the actual linear correlation in interpolated images, given that the neighborhood size K is properly chosen. This is also demonstrated by Figure 1 that depicts a series of predictor coefficients, estimated from a 150 % ($\omega = 2/3$) bilinearly upsampled image. Observe the periodic structure of the estimated coefficients, which (on average) perfectly obey the expected numerical configuration.

Unfortunately, despite being very accurate in principle, TLS has some severe practical limitations in our field of application. The most fundamental drawback lies in the assumption of a vanishing model error. Since TLS assumes homoscedastic error terms, a large model error can lead to incorrect or instable coefficient estimates. This is especially the case for downsampling. However, also an inappropriate predictor size can raise problems (whereas too small neighborhoods are principally more troublesome than too large neighborhoods).[2] While the literature on TLS discusses corrective measures like regularization or scaled error terms [14], we found in our experiments that a weighted least squares approach is generally more adequate for the sole detection of interpolation artifacts.

3.2 A practical approach

The above Sect. 2 mentioned that the correlation between interpolated samples will, depending on the sample's relative position δ_x, gradually vary in the bias towards left and right neighbors. The detection of interpolation artifacts thus not necessarily requires a perfect knowledge of the coefficients $\boldsymbol{\alpha}^{(i)}$. On a coarser level, it may be rather sufficient to analyze the relative influence of left and right neighbors, respectively.

In a simplified 2D linear model, we refrain from making rounding errors explicit and write pixels of the i-th row as a linear combination of their vertical neighbors, superimposed with an additive error term $\boldsymbol{\varepsilon}$,

$$\boldsymbol{r}^{(i)} = \mathbf{R}^{(i)} \cdot \boldsymbol{\beta}^{(i)} + \boldsymbol{\varepsilon}^{(i)} . \tag{5}$$

In the course of this paper, we will employ a weighted least squares (WLS) procedure to estimate the coefficients $\hat{\boldsymbol{\beta}}^{(i)}$, such that

$$\hat{\boldsymbol{r}}^{(i)} = \mathbf{R}^{(i)} \cdot \hat{\boldsymbol{\beta}}^{(i)} \quad \text{and} \quad \left\| \boldsymbol{w}^{(i)} \left(\boldsymbol{r}^{(i)} - \hat{\boldsymbol{r}}^{(i)}\right) \right\| \to \min .$$

The weights $w_j^{(i)}$ are chosen to mitigate the impact of highly textured regions and are computed from the variance σ_j^2 of the pixel's $2K$ vertical neighbors, $w_j^{(i)} \propto 1/(c + \sigma_j^2)$, $c > 0$.

Since it is mainly the relative influence of upper and lower neighbors that is of interest in our setting, we focus our analysis on the coefficients $\beta_{-1}^{(i)}$ and $\beta_1^{(i)}$, which correspond to the direct vertical neighbors. Experiments suggest that the differences

$$d_i = \beta_{-1}^{(i)} - \beta_1^{(i)} \tag{6}$$

are particularly promising to detect traces of resizing.

When analyzing a sequence of predictor coefficients (or their differences), we expect to observe periodic interpolation artifacts if the image was resized. Typically, a detection of such periodicities is carried out in the frequency domain, where the existence of a distinct spectral component is taken as an indicative measure. It is worth mentioning that analyzing the differences d_i instead of the predictor coefficients itself, as a by-product, inherently reduces the influence of undesired low-frequency components, which might stem from the image content. Damping functions [12] or cut-offs [4, 8] are thus typically not necessary.

Realistically, especially downsampling will hardly result in perfectly periodic predictor coefficients $\hat{\boldsymbol{\beta}}^{(i)}$. This means that the corresponding interpolation artifacts can be difficult to detect with traditional spectral methods such as Fourier analysis, as employed for instance in [12, 4, 10, 8]. We will rather turn to a robust spectral estimator that can better deal with "imperfect" signals. More specifically, we employ Ahdesmäki et al.'s spectral density estimator [1], which is based on Spearman's rank correlation coefficient ϱ:

$$S(f) = \sum_{l=-L}^{L} \varrho(l) \exp(-2\pi i f l)$$

[2] Recall that the optimal neighborhood size is unknown in a realistic forensic setting.

Table 1: Detection of downsampling with row/column (RC) predictors and Popescu and Farid's [12] global predictor (GP), respectively. Detection rate at FAR $\leq 1\,\%$ and area under the ROC curve (AUC) for varying scaling factors ω^{-1} and different interpolation kernels*.

ω^{-1}	0.50 RC/GP	0.55 RC/GP	0.60 RC/GP	0.65 RC/GP	0.70 RC/GP	0.75 RC/GP	0.80 RC/GP	0.85 RC/GP	0.90 RC/GP	0.95 RC/GP
detection rate at 1 % false acceptance rate ($TP_{0.01}$)										
Bilinear	0.01 / 0.01	**1.00** / 0.99	**1.00** / 0.41	**0.94** / 0.60	**0.95** / 0.62	**0.75** / 0.58	**1.00** / 0.82	1.00 / 1.00	1.00 / 1.00	1.00 / 1.00
Bicubic	0.01 / 0.01	**0.17** / 0.08	**0.62** / 0.32	**0.78** / 0.66	**0.84** / 0.70	**0.99** / 0.75	**0.99** / 0.82	**0.99** / 0.84	**0.99** / 0.93	**0.99** / 0.97
Lanczos	0.01 / 0.01	**0.14** / 0.05	**0.64** / 0.29	**0.81** / 0.66	**0.87** / 0.70	**0.86** / 0.66	**0.80** / 0.56	**0.73** / 0.48	**0.85** / 0.48	**0.92** / 0.67
Spline	0.01 / 0.01	**0.03** / 0.01	**0.35** / 0.02	**0.61** / 0.12	**0.69** / 0.18	**0.67** / 0.19	**0.59** / 0.12	**0.52** / 0.10	**0.49** / 0.10	**0.38** / 0.05
area under the ROC curve (AUC)										
Bilinear	**0.53** / 0.52	1.00 / 1.00	**1.00** / 0.98	**0.99** / 0.95	**1.00** / 0.97	0.95 / 0.95	**1.00** / 0.99	1.00 / 1.00	1.00 / 1.00	1.00 / 1.00
Bicubic	0.51 / **0.56**	0.68 / **0.81**	0.91 / 0.91	0.96 / 0.96	**0.98** / 0.97	**1.00** / 0.98	**1.00** / 0.99	**1.00** / 0.99	**1.00** / 0.99	1.00 / 1.00
Lanczos	0.48 / **0.58**	0.64 / **0.76**	0.90 / 0.90	0.97 / 0.97	0.97 / 0.97	**0.97** / 0.96	**0.96** / 0.94	**0.95** / 0.92	**0.98** / 0.94	**0.99** / 0.98
Spline	**0.58** / 0.44	**0.60** / 0.46	**0.81** / 0.57	**0.91** / 0.74	**0.93** / 0.80	**0.93** / 0.82	**0.91** / 0.78	**0.91** / 0.75	**0.91** / 0.78	**0.88** / 0.75

*corresponding ImageMagick `-filter` setting: `Triangle`, `Catrom`, `Lanczos`, `Cubic`.

We conclude this section with a remark on the comparably moderate computational resources that a series of row predictors requires. While the number of regressions obviously grows with the number of rows in the image under investigation, it is important to mention that each regression entails by far fewer samples than Popescu and Farid's iterative global estimator [12]. More specifically, a $N \times M$ image involves N independent row predictors with each $2KM$ samples, whereas one iteration of the global $(2K+1) \times (2K+1)$ predictor is based on $4K(K+1)MN$ samples. This is of particular interest in the analysis of large(r) images, where the latter can quickly impose tremendous memory requirements. Compared to detectors based on fixed linear kernels [7], regression-based methods will always be computationally more demanding. However, as we shall see in the following section, this is leveled out by an increased reliability.

4. EXPERIMENTAL RESULTS

For a demonstration of the general applicability of row (and column) predictors in the detection of resized images, we make use of (a part of) the 'Dresden image database' [5]. More specifically, approximately 1100 never-compressed images, stemming from five different digital camera models, were resized by various amounts using ImageMagick's `convert` (version `6.6.2-8`) with altogether four different interpolation kernels. The interpolation kernels were chosen to cover the range from rather simple methods, such as bilinear interpolation (`-filter Triangle`), to more sophisticated algorithms, such as bicubic (`-filter Catrom`), windowed sinc (`-filter Lanczos`), or cubic spline (`-filter Cubic`) interpolation. The analysis of interpolation artifacts was then conducted using the center 512×512 region of the green channel of the scaled images.

Our detector employs weighted least squares predictors with a neighborhood size $K = 3$, weights w_j with $c = 10^{-5}$ (for pixel intensities measured in the range $[0, 1]$), and relies on the decision criterion ρ,

$$\rho = \max_f S(f) \,/\, \operatorname*{median}_f S(f)\,,$$

with $S(f)$ being the spectral density estimate [1] of the differences d_i, cf. Sect. 3.2. We expect interpolated images to exhibit a distinct peak, and thus images with $\rho > T$, with T being a predefined threshold, are flagged as resized. To make use of row *and* column predictors, the detector finally decides based on the maximum decision criterion, $\max(\rho_{\text{row}}, \rho_{\text{column}})$. The detection thresholds for a given false acceptance rate (FAR) were obtained by running the detector on all original images in the test database.

For comparison, all images were additionally analyzed with Popescu and Farid's detector [12], which is usually considered to be the most reliable detector, especially for downsampling [7]. The neighborhood of the global predictor was set to $K = 2$, i. e., a window of size 5×5, and the detector performed an exhaustive search over 200 synthetic downsampling maps with $\omega^{-1} \in [0.5, 1)$, sampled in equidistant steps of 1/400.

Since it is generally accepted that upsampling is very well detectable, we will center the following discussion on our results from downsized images. Table 1 gives an overview by reporting detection rates (true positives) at a false positive rate of $\leq 1\,\%$, $TP_{0.01}$, together with AUC (area under the ROC curve) values for both the row/column predictors (RC) and global predictor (GP). For better visualization, parts of the results, namely graphs for bilinear and bicubic interpolation, are also depicted in Fig. 2.

Before we delve into a more detailed analysis of the numerical results, the so far unmatched size of our employed test database justifies a short note to support general findings from the literature. To be more specific, bilinear interpolation, due to its smaller support, is better detectable than more sophisticated interpolation algorithms, even though we can spot some exceptions for particular scaling factors. Also, stronger downsampling typically leaves behind less detectable traces.

What is more interesting in our setting is the relatively good performance of the row/column predictors. As indicated by Tab. 1 and Fig. 2 (a), row/column predictors outperform Popescu and Farid's global predictor in particular in the important case of small false positive rates. This becomes most visible for downsampling with more sophisticated interpolation algorithms, even though the overall results may seem not very convincing when compared to figures reported in prior work. We should however keep in mind, that ImageMagick automatically applies some sort of

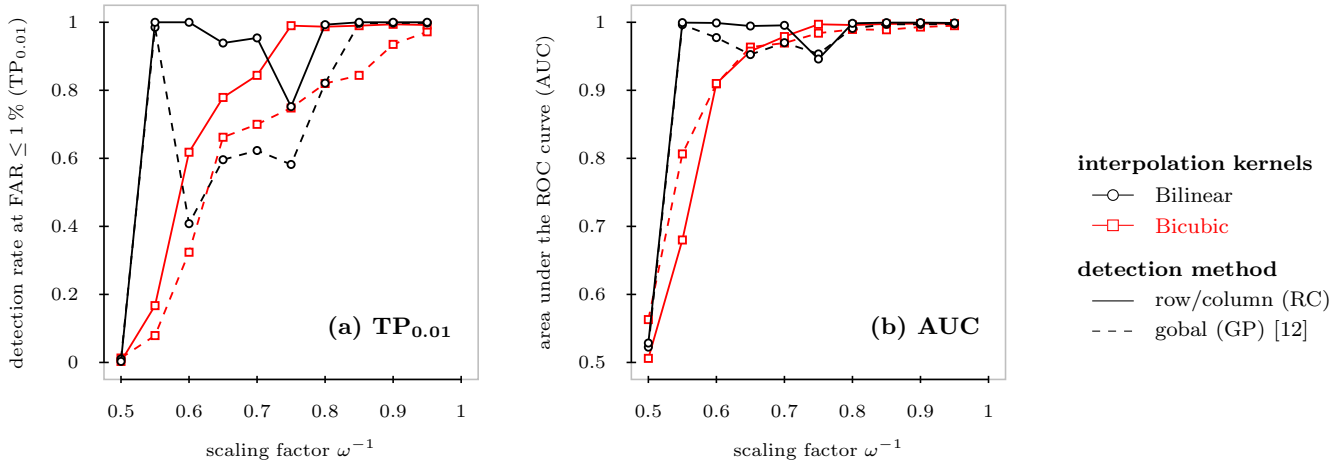

Figure 2: Detection of downsampling with row/column (RC) predictors and Popescu and Farid's [12] global predictor (GP), respectively. (a) Detection rate at FAR $\leq 1\,\%$, and (b) area under the ROC curve (AUC) for varying scaling factors ω^{-1}. Bilinear and bicubic interpolation.

post-processing (for instance anti-aliasing and smoothing) to create visually more pleasing images. While this admittedly distracts the pure scientific view on the detectability of interpolation artifacts to a certain degree, we nevertheless deliberately decided to use 'real-world' software in our experiments to gain more realistic insights.

Figure 2 (a) suggests that row/column predictors are particularly favorable when striving for low false positive rates. The gain in overall performance in terms of AUC values however strongly depends on the interpolation kernel. On the one hand, as depicted in Fig. 2 (b), the differences for bilinear and bicubic interpolation are rather negligible. The detection of Lanczos and cubic spline interpolation, on the other hand, clearly benefits from an analysis with row/column predictors, cf. Tab. 1.

Since $\mathrm{TP}_{0.01}$ and AUC measures provide only a limited view on the actual detection performance, accompanying ROC curves for downsampling to 80 % of the original image size ($\omega^{-1} = 0.8$) are depicted in Fig. 3. We found this scaling factor to be a good representative of the detectors' typical characteristics. Note how the relative overall performance gain of the row/column (RC) predictors over Popescu and Farid's global predictor (GP) increases from bilinear to cubic spline interpolation.

In general, our experimental results suggest that row and column predictors, combined with robust spectral estimators, form a valuable alternative in the analysis and detection of resized images. While the reported results of the row/column predictors exclusively stem from an analysis with $K = 3$, it is worth mentioning that complementary experiments with larger neighborhoods showed no noticeable improvements. Predictors that include fewer neighbors ($K \leq 2$) may—in particular for stronger downsampling—not always be as reliable.

5. CONCLUDING REMARKS

In this paper, we have explored how the characteristic structure of resized (and thus interpolated) images can be exploited to model the typical linear correlations between neighboring pixels more explicitly. While existing methods measure interpolation artifacts in some sort of residue signal, obtained by linear filtering with a global kernel, we have shown that the specific periodicities can also be detected in a series of tailored row and column predictors. The underlying key observation was that, in resized images, all pixels within one row are correlated with their vertical neighbors in exactly the same way (and correspondingly for columns and their horizontal neighbors). Tests on a large database of images demonstrated that an explicit model of linear dependencies in scaled images forms a promising forensic approach to detect resizing.

Apart from the discussed application in the detection of interpolation artifacts, row and column predictors offer additional valuable information that can be useful in a forensic setting. Future research will investigate how the estimated predictor coefficients can be exploited to distinguish upsampled from downsampled images. It is well-known that a whole set of geometric transformations leads to the same spectral artifacts in the analysis of interpolated images, rendering a distinction impossible [12, 7]. At the same time, however, different scaling factors result in varying predictor coefficients, which gives rise to a possible disambiguation.

A further interesting application may be found in the steganalysis of (down)scaled images, where an explicit model of linear dependencies could help to better describe and estimate stego noise, cf. for instance [6].

As to the limitations, we have to point out that row and column predictors are limited to the analysis of resized images. Extensions to more general geometric transformations are not straight-forward (or even impossible). A probably more severe, but also more general problem lies in the reliable detection of downsampling with more or less sophisticated kernels beyond bilinear interpolation. While the reported results show clear improvements over prior art, the achievable detection rates admittedly still preclude an application in practically relevant scenarios. It remains an open question whether targeted detectors based on periodic interpolation artifacts will ever reach a satisfactory level in this regard, or whether we will ultimately end up with a classifier-based approach that relies on tens or hundreds of features [9, as a first step in this direction].

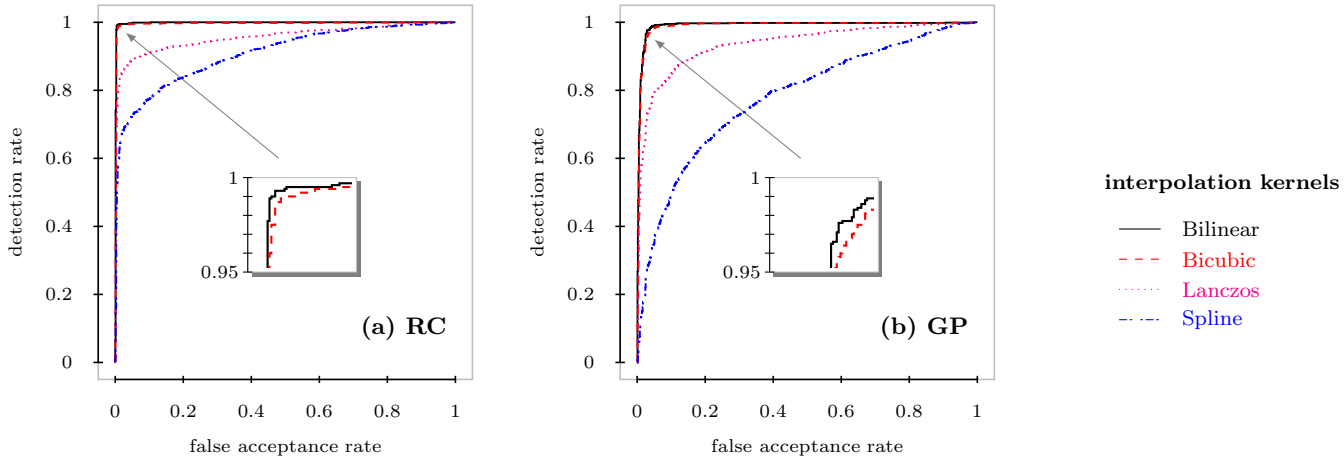

Figure 3: ROC curves for the detection of downsampling with (a) row/column predictors and (b) Popescu and Farid's global predictor [12]. Resizing to 80 % of the original image size.

Acknowledgements

The author gratefully receives a doctorate scholarship from Deutsche Telekom Stiftung, Bonn.

6. REFERENCES

[1] M. Ahdesmäki, H. Lähdesmäki, R. Pearson, H. Huttunen, and O. Yli-Harja. Robust detection of periodic time series measured from biological systems. *BMC Bioinformatics*, 6(117), 2005.

[2] R. Böhme, F. Freiling, T. Gloe, and M. Kirchner. Multimedia forensics is not computer forensics. In Z. J. Geradts, K. Y. Franke, and C. J. Veenman, editors, *Computational Forensics, Third International Workshop, IWCF 2009*, LNCS 5718, pages 90–103, Berlin, Heidelberg, 2009. Springer Verlag.

[3] H. Farid. Image forgery detection. *IEEE Signal Processing Magazine*, 26(2):16–25, 2009.

[4] A. C. Gallagher. Detection of linear and cubic interpolation in JPEG compressed images. In *Second Canadian Conference on Computer and Robot Vision*, pages 65–72, 2005.

[5] T. Gloe and R. Böhme. The Dresden Image Database for benchmarking digital image forensics. In *25th Symposium on Applied Computing (ACM SAC 2010)*, pages 1585–1591, 2010.

[6] A. D. Ker and R. Böhme. Revisiting weighted stego-image steganalysis. In E. J. Delp, P. W. Wong, J. Dittmann, and N. Memon, editors, *SPIE-IS&T Electronic Imaging: Security, Forensics, Steganography, and Watermarking of Multimedia Contents X*, volume 6819, 681905, 2008.

[7] M. Kirchner. Fast and reliable resampling detection by spectral analysis of fixed linear predictor residue. In *Multimedia and Security Workshop 2008*, pages 11–20, New York, NY, USA, 2008. ACM Press.

[8] M. Kirchner and T. Gloe. On resampling detection in re-compressed images. In *First IEEE International Workshop on Information Forensics and Security*, pages 21–25, 2009.

[9] Q. Liu and A. H. Sung. A new approach for JPEG resize and image splicing detection. In *First ACM Workshop on Multimedia in Forensics*, pages 43–48, New York, NY, USA, 2009. ACM Press.

[10] B. Mahdian and S. Saic. Blind authentication using periodic properties of interpolation. *IEEE Transactions on Information Forensics and Security*, 3(3):529–538, 2008.

[11] A. C. Popescu and H. Farid. Statistical tools for digital forensics. In J. Fridrich, editor, *Information Hiding. 6th International Workshop*, LNCS 3200, pages 128–147, Berlin, Heidelberg, 2004. Springer Verlag.

[12] A. C. Popescu and H. Farid. Exposing digital forgeries by detecting traces of re-sampling. *IEEE Transactions on Signal Processing*, 53(2):758–767, 2005.

[13] H. T. Sencar and N. Memon. Overview of state-of-the-art in digital image forensics. In B. B. Bhattacharya, S. Sur-Kolay, S. C. Nandy, and A. Bagchi, editors, *Algorithms, Architectures and Information Systems Security*, volume 3 of *Statistical Science and Interdisciplinary Research*, chapter 15, pages 325–348. World Scientific Press, 2008.

[14] S. Van Huffel and J. Vandewalle. *The Total Least Squares Problem: Computational Aspects and Analysis*. Society for Industrial and Applied Mathematics, 1991.