

“STEGANALYSIS IN TECHNICOLOR”

BOOSTING WS DETECTION OF STEGO IMAGES FROM CFA-INTERPOLATED COVERS

Matthias Kirchner and Rainer Böhme

University of Münster, Dept. of Information Systems, Leonardo-Campus 3, 48149 Münster, Germany

ABSTRACT

Steganographic security in empirical covers is best understood for grayscale images. However, the world, and almost all digital images of it, are more colorful. This paper extends the weighted stego-image (WS) steganalysis method to detect stego images produced from covers that exhibit traces of color filter array (CFA) interpolation, which is common for images acquired with digital cameras. The approach combines techniques of CFA forensics with state-of-the-art WS steganalysis. Empirical results from large datasets indicate significant increases in detection performance, in particular for small payloads. This specific weakness of color covers calls into question the common assumption that grayscale image steganography generalizes to color images by treating each chroma channel independently.

Index Terms— Steganalysis, color filter array, image forensics

1. INTRODUCTION

A steganographic system is secure if stego objects are *plausible*; that is, their distribution is *indistinguishable* by an adversary from the distribution of cover objects transmitted on a channel. Most published and almost all well-understood steganographic embedding functions are only secure—if at all—against adversaries suffering from achromatopsia: they are designed for grayscale images only. Few things demonstrate more impressively how detached steganography research is from the real world [1]. Arguably, a scheme offering some security for monochrome images can be generalized to color images by applying the same embedding function on each chrominance channel of an image. Early implementations of embedding functions take this approach (e. g., F5 and OutGuess). However, the security implications are not well understood and devastating surprises may loom.

This paper sets out to show that known steganalysis techniques can be substantially improved by considering the provenance of color information in suspect images. More specifically, we replace the cover predictor in weighted stego-image (WS) steganalysis [2, 3] by position-specific predictors to account for differences in the local predictability of pixels depending on their position in the color filter array (CFA) interpolation. CFA interpolation leads to systematic and highly variable predictability of cover pixels. It is virtually unavoidable when acquiring plausible covers with commercial digital cameras because CFAs are physical parts built into the optical system [4]. Our simple and run-time neutral extension yields significant increases in detectability, which translate to a security weakness whenever embedding functions designed and evaluated for grayscale images are naively generalized to color covers.

We are not aware of any other detector explicitly using color information for this purpose. A proposal to count neighbors in the color cube to detect additive noise steganography in RGB images [5] later turned out to depend on artifacts of color subsampling in JPEG

Table 1. Relevant linear 3×3 filter kernels

\mathcal{F}_{4N}	\mathcal{F}_{4D}	\mathcal{F}_{2H}	\mathcal{F}_{2V}
$\begin{bmatrix} 0 & \frac{1}{4} & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & \frac{1}{4} & 0 \end{bmatrix}$	$\begin{bmatrix} \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & 0 & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} \end{bmatrix}$	$\begin{bmatrix} 0 & 0 & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} \\ 0 & 0 & 0 \end{bmatrix}$	$\begin{bmatrix} 0 & \frac{1}{2} & 0 \\ 0 & 0 & 0 \\ 0 & \frac{1}{2} & 0 \end{bmatrix}$
\mathcal{F}_{KBS}	\mathcal{F}_{LS8}	\mathcal{F}_{green}	\mathcal{F}_{red}
$\begin{bmatrix} -\frac{1}{4} & \frac{1}{2} & -\frac{1}{4} \\ \frac{1}{2} & 0 & \frac{1}{2} \\ -\frac{1}{4} & \frac{1}{2} & -\frac{1}{4} \end{bmatrix}$	$\begin{bmatrix} b & a & b \\ a & 0 & a \\ b & a & b \end{bmatrix}$	$\begin{bmatrix} 0 & \frac{1}{4} & 0 \\ \frac{1}{4} & 1 & \frac{1}{4} \\ 0 & \frac{1}{4} & 0 \end{bmatrix}$	$\begin{bmatrix} \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{2} & 1 & \frac{1}{2} \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \end{bmatrix}$

pre-compressed covers [6]. It can therefore be seen as a predecessor of this work, even if not designed with this goal in mind.

As to the organization of this paper, Section 2 defines notational conventions and recalls the essentials of WS steganalysis and CFA interpolation. Section 3 combines the two to a more powerful detector, which is then evaluated in Section 4. The final Section 5 concludes with implications for future research.

2. PRELIMINARIES

2.1. Notation and Definitions

Boldface symbols denote vectors and matrices. We represent individual channels of a n -pixel RGB color image, \mathbf{x} , as vectorized integer intensity lattices, $\mathbf{x} = (x_1, \dots, x_{3n}) = (\mathbf{x}_{\{red\}}, \mathbf{x}_{\{green\}}, \mathbf{x}_{\{blue\}}) \in \mathbb{Z}^{3n}$. If not stated otherwise, we omit the channel subscript and examine color channels independently. Symbol $x_i^{(p)}$ denotes an intensity value after steganographic embedding with net embedding rate $p \in [0, 1]$, i. e., $\mathbf{x}^{(0)}$ is a cover image. We write $\mathcal{F}(\mathbf{x})$ to refer to a linearly filtered version of intensity lattice \mathbf{x} . Table 1 gives an overview of filter masks we will employ throughout the paper. The first five predict pixel values from their local spatial neighborhood. \mathcal{F}_{LS8} has the same purpose, yet instead of fixed coefficients, a and b are found adaptively in a least squares (LS) procedure by minimizing the L^2 distance between a given image and predicted values. Filters \mathcal{F}_{green} and \mathcal{F}_{red} will be useful to formalize green and red channel bilinear CFA interpolation, respectively.

2.2. WS Steganalysis in a Nutshell

Fridrich and Goljan’s quantitative weighted stego-image (WS) steganalysis method [2] estimates the embedding rate \hat{p} of uniform least significant bit (LSB) replacement embedding in (grayscale) intensity

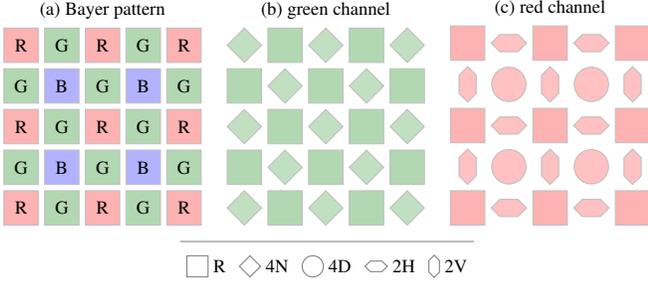


Fig. 1. Spatial neighborhood types in CFA-interpolated images; raw (R), amidst four adjacent raw pixels (4N), four diagonal raw neighbors (4D), two horizontal (2H) or vertical (2V) raw neighbors. The legend symbols reappear as point marks in Figures 3 and 4.

images. In its most basic version, the estimator takes the form

$$\hat{p} = \frac{2}{n} \sum_{i=1}^n (-1)^{x_i^{(p)}} \left(x_i^{(p)} - \hat{x}_i^{(0)} \right). \quad (1)$$

The first factor in the summation determines the sign of the pixel’s contribution to the estimate based on its LSB. Vector $\hat{x}^{(0)}$ is an estimate of the cover pixel, obtained by linear filtering of the stego image, $\hat{x}^{(0)} = \mathcal{F}(x^{(p)})$. Equation (1) is a consistent estimator of the embedding rate p if $\hat{x}_i^{(0)}$ does not depend on $x_i^{(p)}$. Optional local weights w_i may be used to account for variation in the local predictability of $x_i^{(0)}$. The original WS proposal uses the mean of the surrounding four stego pixels to predict the cover pixel, $\mathcal{F} = \mathcal{F}_{4N}$. Subsequent works have proposed filters of the form \mathcal{F}_{KB8} or adaptive filters \mathcal{F}_{LS8} to steganalyze never-compressed grayscale images [3].

A convenient summary measure for the detection performance is the absolute estimation error $|\hat{p} - p|$, which can be averaged over multiple covers to compute the mean absolute error (MAE) metric. Lower values of MAE imply better detection performance. By this measure, enhanced WS detectors are reported to be the most sensitive targeted steganalyzers of LSB replacement steganography [3, 7].

2.3. CFA Interpolation in a Nutshell

Most digital cameras combine a single sensor with a color filter array (CFA), i. e., individual sensor elements capture specific color information only. A full-color image is obtained through interpolation from surrounding samples of the raw signal in a so-called demosaicing procedure [4]. The most common CFA layout is the family of Bayer patterns [8], which implies that at least two thirds of all intensity values in an RGB image are interpolated (cf. Fig. 1-a). Because of the Bayer patterns’ periodic structure, this procedure is likely to leave local correlation artifacts between pixels. Their presence, strength and form can be measured particularly well in high-pass filtered versions of the respective color channels [9, 10, 11]. Different combinations of CFA layout and interpolation function yield different artifacts, which gives rise to forensic image source identification algorithms [10, 11]. Tests for a consistent presence of CFA artifacts may expose image manipulations [9, 12].

3. PROPOSED METHOD

To facilitate the following analyses, we make the simplifying (but not limiting) assumption of plain bilinear CFA interpolation. The respective green and red channel interpolation filters are given in

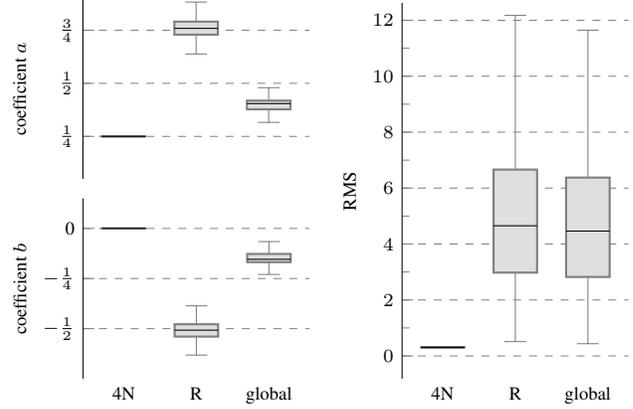


Fig. 2. Impact of CFA neighborhood relations on estimated optimal \mathcal{F}_{LS8} filter coefficients (left) and RMS difference between image and prediction (right). Box plots from 7,408 BOSSBase images after bilinear green channel interpolation.

Tab. 1. The blue channel filter is equivalent to \mathcal{F}_{red} , shifted by one pixel horizontally and vertically.

3.1. Improved Local Predictions in CFA-Interpolated Covers

For bilinear interpolation, the specific form of CFA artifacts intuitively follows from the different spatial *neighborhood types* of a Bayer pattern (cf. Fig. 1). Specifically, there exist two classes of green color channel pixels. Raw pixels, $\mathbf{x}_{\{R\}}$, are located at sites where the CFA naturally has a green element. Interpolated pixels, $\mathbf{x}_{\{4N\}}$, have four raw pixels as direct neighbors. Interpolation filter \mathcal{F}_{green} dictates that pixels of this type are equal to the mean of their non-interpolated 4-neighborhood (plus a rounding error ϵ),

$$(\mathbf{x} - \mathcal{F}_{4N}(\mathbf{x}))_{\{4N\}} = \epsilon. \quad (2)$$

As for the red and blue color channels, we can distinguish between three different types of interpolated pixels, namely $\mathbf{x}_{\{2H\}}$, $\mathbf{x}_{\{2V\}}$, and $\mathbf{x}_{\{4D\}}$. Strong neighborhood correlations, as in Eq. (2), are obtained by inserting the appropriate filter kernels from Tab. 1. The predictability of raw pixels is generally lower because they depend to a larger extent on the image content.

Figure 2 exemplarily shows the impact of different CFA neighborhood types on the predictability of pixels after green channel Bayer pattern demosaicing of 7,408 BOSSBase images¹ [13]. The box plots in the left panel of the figure summarize the distributions of estimated optimal \mathcal{F}_{LS8} coefficients a and b for raw (R) and interpolated (4N) green channel pixels. For comparison, we also report the respective “global” coefficients, which we estimate from all pixels disregarding their CFA neighborhood type. The right panel displays corresponding box plots of per-image root mean square differences between actual pixel intensities and predicted values. The empirical distributions emphasize the high predictability of interpolated pixels. For this type, the estimated coefficients adhere to Eq. (2) with negligible deviations ($a = 1/4, b = 0$) and yield very low prediction errors (median RMS: 0.3). Coefficient estimates for non-interpolated pixels differ considerably (median a : 0.76, median b : -0.51) and are subject to content-dependent variation. The resulting prediction errors indicate a substantially lower predictability (median RMS:

¹A precise description of the dataset is given in Sect. 4.1.

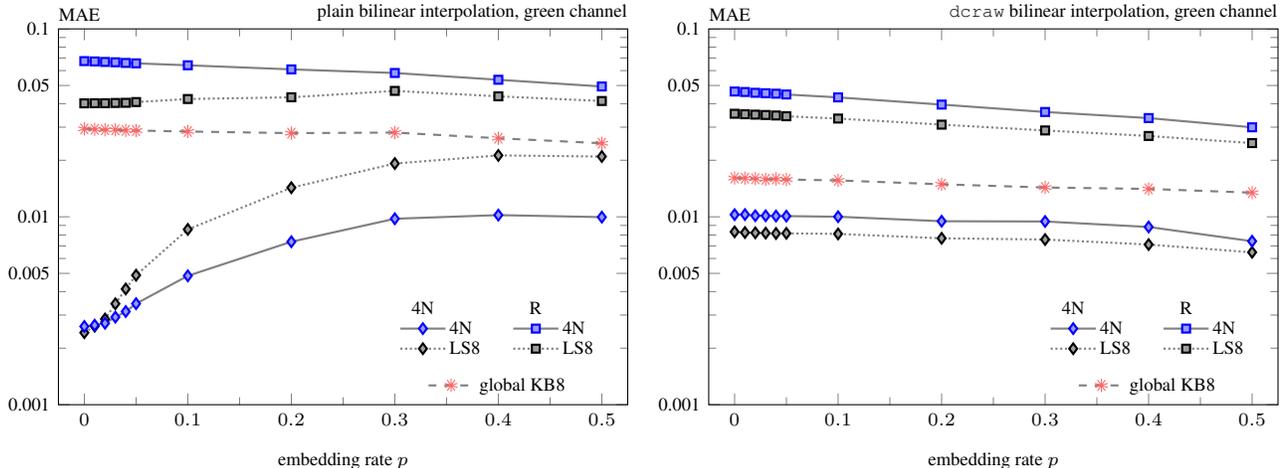


Fig. 3. Steganalysis results for different CFA neighborhood types as a function of embedding rate; left: 7,408 BOSS images; right: 3,316 Dresden Image Database color image blocks (Nikon D70); green channel bilinear CFA interpolation.

4.65). Global estimation results are a mixture of the two individual types, with an overall slightly better pixel predictability than for raw pixels alone. Interestingly, globally optimal filter coefficients (median a : 0.40, median b : -0.15) moderate the drastic differences in the predictability of raw and interpolated pixels mostly at the expense of larger prediction errors for the latter type. Specifically, we observe a median RMS of 2.31 when applying global coefficients to predict interpolated pixels only, whereas the median RMS for raw pixels is 5.85. The results for the red and blue channel are very similar and are thus omitted here for the sake of brevity.

3.2. WS Steganalysis with Type-Specific Cover Predictors

The performance of the estimator \hat{p} in Eq. (1) strongly depends on the quality of the cover image estimate $\hat{x}^{(0)}$ [3]. Hence, we deem an explicit utilization of the specific CFA neighborhood relations beneficial to steganalyze (individual channels of) color images. This holds in particular for the types of interpolated pixels which have tailored predictors as in Eq. (2). We thus replace the global cover predictor \mathcal{F} with a type-specific version \mathcal{F}_C , $C \in \{4N, 4D, 2H, 2V, R\}$, depending on each individual pixel’s CFA position. Type-specific WS estimates of the unknown embedding rate are then given by

$$\hat{p}_C = \frac{2}{|\{C\}|} \sum_{\{i \in C\}} (-1)^{x_i^{(p)}} \left(x_i^{(p)} - \mathcal{F}_C(x^{(p)})_i \right). \quad (3)$$

Optional local weights can be applied. Aggregating type-specific estimates \hat{p}_C to a combined estimate \hat{p} is equivalent to assigning weights.

4. EXPERIMENTAL VALIDATION

4.1. Data and Setup

We use the BOSSBase [13] and a subset of the Dresden Image Database [14] for our experiments. The former contains 10,000 grayscale images of size 512×512 , downsized from full-resolution digital camera images. This substantial shrinking should remove all genuine CFA artifacts. Then, we synthesize RGB cover images with ideal CFA artifacts by sampling the grayscale images onto the Bayer

grid in Fig. 1-a and demosaic them with plain bilinear CFA interpolation. We further use the tool `drcraw` with bilinear interpolation to generate 3,400 color images from Nikon D70 raw images in the Dresden Image Database. This tool produces more realistic output than plain bilinear interpolation as it applies more of the color processing pipeline, such as white balancing. The size of the `drcraw` images is also 512×512 , obtained by randomly cropping five non-overlapping blocks from each demosaiced full-resolution image in landscape format. The cropping positions take the 2×2 CFA periodicity into account, so that all blocks share the same CFA layout. A third set of covers has been produced using Adobe Lightroom instead of `drcraw`. This tool does not use bilinear interpolation, but demosaics raw images with a sophisticated (proprietary) content-adaptive function.

For each cover, steganographic embedding with embedding rates $p \in \{0.01, 0.02, 0.03, 0.04, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$ is simulated by flipping the least significant bits (LSBs) of randomly chosen sets of $np/2$ pixels per color channel. We report steganalysis results for individual color channels to isolate the effects of different CFA neighborhoods. In all our analyses, we follow the practice to exclude covers with more than 5% flat blocks (of size 3×3), because WS is known to accumulate bias from flat regions [2], and we do not consider bias correction here. This selection explains the odd numbers of images in the reported results. (Cautious steganographers should never embed into covers with flat areas anyways.)

4.2. Steganalysis Results

4.2.1. Green channel

Figure 3 reports detection performance measured by MAE as a function of the embedding rate. The “global” KB8 detector represents the state of the art and serves as benchmark. If the analysis is constrained on interpolated pixels in the green channel using the type-specific predictors, detection performance increases by up to an order of magnitude for small embedding rates (i. e., the most relevant scenario). The advantage of the proposed method is also visible for the `drcraw` images, but substantially smaller. This can be explained by the increasing model mismatch between estimates from predictors based on a image model that assumes plain interpolation and a more complex post-processing pipeline. Observe that the relative performance of

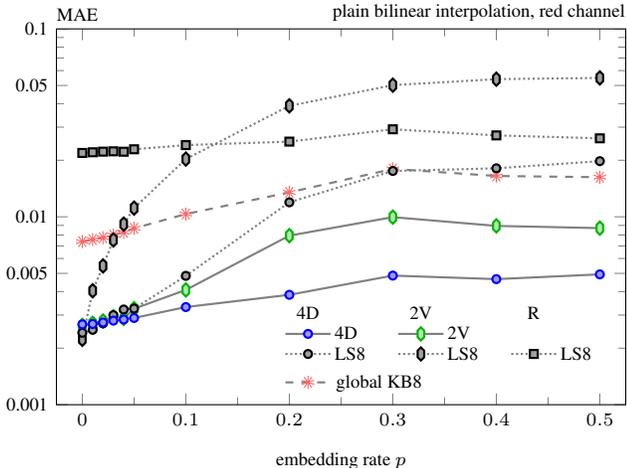


Fig. 4. Steganalysis results for different CFA neighborhood types as a function of embedding rate; 6,127 BOSS images; red channel plain bilinear CFA interpolation.

fixed and adaptive predictors changes between plain bilinear interpolation and `dcraw`. To a certain degree, the adaptive predictors can better adjust to slight model mismatch. Quite expectedly, the estimation errors are higher when only raw pixels are evaluated.

Table 2 summarizes performance indicators for the binary hypothesis test $p > 0$ using selected predictors and CFA-interpolated covers with and without embedding at $p = 0.01$. These figures also show the performance boost for plain bilinear interpolation, which is attenuated but still significant for the `dcraw` images. Only Lightroom causes a model mismatch large enough to let the KB8 predictor gain advantage over type-specific predictors. We speculate that part of this disadvantage can be compensated with larger kernel sizes and content as well as type-adaptive predictors.

4.2.2. Red and blue channels

Figure 4 shows the corresponding results for the red channel. We also find an advantage of the proposed method over the KB8 state-of-the-art detector. The 4D predictor performs best. This is unsurprising as the pixels it predicts exhibit the strongest linear dependence. More surprisingly, the predictors 2V and 2H (not shown) outperform the KB8 although they only take into account the information from two neighbors, which makes them more sensitive to bias.

Due to the shared filter configurations, the blue channel results are similar to the red channel.

4.3. Practical Considerations

Evaluating Eq. (3) with pre-set filter coefficients requires prior knowledge of the underlying CFA layout. If this information is not directly available, it can be inferred from the suspect image [15]. This resembles the “forensics-aided” approach to steganalyze heterogeneous material: in a first step, a forensic classifier determines the cover source, then, a tailored method is picked from a bank of steganalyzers to output the final estimate or decision [16]. Alternatively, it is possible to find the optimal filter coefficients adaptively using an LS procedure with two sets of coefficients for the green channel, and four each for the red and blue channels. The latter approach increases the running time by a constant factor, but it remains linear in n .

Table 2. Equal error rate (EER) and false positive rate at 50% detection rate (FP_{50}) for binary steganalysis decisions on green channel CFA-interpolated covers.

\mathcal{F}	$(p = 0.01)$		bilinear		adaptive	
			plain	<code>dcraw</code>	Lightroom	
			$N=7,408$	$N=3,316$	$N=3,166$	
	FP_{50}	EER	FP_{50}	EER	FP_{50}	EER
Standard WS (KB8)						
\mathcal{F}_{KB8}	0.35	0.41	0.23	0.33	0.08	0.19
\mathcal{F}_{LS8}	0.38	0.43	0.26	0.34	0.09	0.20
Proposed CFA-WS (4N)						
\mathcal{F}_{4N}	0.01	0.05	0.15	0.20	0.19	0.30
\mathcal{F}_{LS8}	0.01	0.04	0.10	0.16	0.11	0.23

5. CONCLUSION AND OUTLOOK

We propose an improved variant of WS steganalysis optimized for stego images from CFA-interpolated covers (shorthand: CFA-WS). Experimental results show substantial performance boosts for small embedding rates as long as sufficient information about the demosaicing function is available or can be inferred from the suspect image. In numbers: a steganalyst who faces 35% false positives at 50% detection rate with conventional targeted detectors can reduce her false positives to 1% by using the proposed method.

The analysis was intentionally confined to LSB replacement because it is best understood and proven efficient closed-form detectors are readily available [17]. We also refrained from aggregating the evidence extracted from each color channel as this is related to assigning weights w_i in weighted WS steganalysis (see Sections 2.2 and 3.2), a topic that is best studied independently [18]. Next steps include other embedding operations in the spatial and possibly transformed domain, as well as considering more general preprocessing chains, such as interpolation after resizing.

Our findings illustrate how urgent it is to fill the gap of rigorous research on color image steganography.² Color-sensitive steganalysis is only a first step that can define a benchmark. And our results must be interpreted as lower bound for the additional insecurity of embedding in color images because our detector exploits only dependencies *within* a color channel due to color interpolation at cover generation. Dependencies *between* color channels remain to be explored for steganography, suggesting detection strategies that verify the consistency of CFA artifacts throughout a suspect image [9] and embedding strategies that integrate methods to synthesize a plausible CFA pattern in stego images [20]. Unless those techniques are well understood, steganographers should better stay away from color images (and hence from image steganography at all unless they find a channel where grayscale images are plausible).

*

The title is an allusion to the Coldplay song “Life in Technicolor”. Similarities with existing companies and trademarks are accidental.

This research was funded by Deutsche Forschungsgemeinschaft (DFG) under grant “Sichere adaptive Steganographie”.

²MPSteg-color is a commendable exception [19].

6. REFERENCES

- [1] Andrew D. Ker, Patrick Bas, Rainer Böhme, Rémi Cogranne, Scott Craver, Tomáš Filler, Jessica Fridrich, and Tomáš Pevný, “Moving steganography and steganalysis from the laboratory into the real world,” in *1st ACM Workshop on Information Hiding and Multimedia Security (IH&MMSec '13)*, Montpellier, France, June 2013, pp. 45–58, ACM Press.
- [2] Jessica Fridrich and Miroslav Goljan, “On estimation of secret message length in LSB steganography in spatial domain,” in *Proceedings of SPIE-IS&T Electronic Imaging: Security, Steganography and Watermarking of Multimedia Contents VI*, Edward J. Delp and Ping Wah Wong, Eds., San Jose, CA, 2004, SPIE.
- [3] Andrew D. Ker and Rainer Böhme, “Revisiting weighted stego-image steganalysis,” in *Proceedings of SPIE-IS&T Electronic Imaging: Security, Forensics, Steganography, and Watermarking of Multimedia Contents X*, Edward J. Delp, Ping Wah Wong, Nasir D. Memon, and Jana Dittmann, Eds., San Jose, CA, 2008, vol. 6819, 681905, SPIE.
- [4] Rajeev Ramanath, Wesley E. Snyder, Youngjun Yoo, and Mark S. Drew, “Color image processing pipeline,” *IEEE Signal Processing Magazine*, vol. 34, no. 1, pp. 34–43, 2005.
- [5] Andreas Westfeld, “Detecting low embedding rates,” in *Information Hiding, 5th International Workshop 2002*, Fabien A. P. Petitcolas, Ed., Berlin Heidelberg, 2003, vol. 2578 of *Lecture Notes in Computer Science*, pp. 324–339, Springer.
- [6] Andrew D. Ker, “Resampling and the detection of LSB matching in colour bitmaps,” in *Proceedings of SPIE-IS&T Electronic Imaging: Security, Steganography and Watermarking of Multimedia Contents VII*, Edward J. Delp and Ping W. Wong, Eds., San Jose, CA, January, 16–20 2005, vol. 5681, pp. 1–15, SPIE.
- [7] Rainer Böhme, “Weighted stego-image steganalysis for JPEG covers,” in *Information Hiding*, K. Solanki, K. Sullivan, and U. Madhow, Eds., Berlin Heidelberg, 2008, vol. 5284 of *Lecture Notes in Computer Science*, pp. 178–194, Springer.
- [8] B. E. Bayer, “Color imaging array,” US Patent, 3 971 065, 1976.
- [9] Alin C. Popescu and Hany Farid, “Exposing digital forgeries in color filter array interpolated images,” *IEEE Transactions on Signal Processing*, vol. 53, no. 10, pp. 3948–3959, 2005.
- [10] Andrew C. Gallagher and Tsu-Han Chen, “Image authentication by detecting traces of demosaicing,” in *IEEE Workitorial on Vision of the Unseen (in conjunction with CVPR)*, 2008.
- [11] Hong Cao and Alex C. Kot, “Accurate detection of demosaicing regularity for digital image forensics,” *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 4, pp. 899–910, 2009.
- [12] Ashwin Swaminathan, Min Wu, and K. J. Ray Liu, “Digital image forensics via intrinsic fingerprints,” *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 1, pp. 101–117, 2008.
- [13] Patrick Bas, Tomáš Filler, and Tomáš Pevný, “Break our steganographic system—the ins and outs of organizing BOSS,” in *Information Hiding, 13th International Conference, IH 2011*, Tomáš Filler, Tomáš Pevný, Scott Craver, and Andrew Ker, Eds. 2011, vol. 6958 of *Lecture Notes in Computer Science*, pp. 59–70, Springer.
- [14] Thomas Gloe and Rainer Böhme, “The Dresden Image Database for benchmarking digital image forensics,” *Journal of Digital Forensic Practice*, vol. 3, no. 2–4, pp. 150–159, 2010.
- [15] Matthias Kirchner, “Efficient estimation of CFA pattern configuration in digital camera images,” in *Proceedings of SPIE-IS&T Electronic Imaging: Media Forensics and Security II*, Nasir D. Memon, Jana Dittmann, Adnan M. Alattar, and Edward J. Delp, Eds. 2010, vol. 7541, 7541111, SPIE.
- [16] Mauro Barni, Giacomo Cancelli, and Annalisa Esposito, “Forensics aided steganalysis of heterogeneous images,” in *Proc. of IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. 2010, pp. 1690–1693, IEEE.
- [17] Lionel Fillatre, “Adaptive steganalysis of least significant bit replacement in grayscale natural images,” *IEEE Transactions on Signal Processing*, vol. 60, no. 2, pp. 556–569, 2012.
- [18] Pascal Schöttle, Stefan Korff, and Rainer Böhme, “Weighted stego-image steganalysis for naive content-adaptive embedding,” in *IEEE International Workshop on Information Forensics and Security (WIFS)*, 2012, pp. 193–198, IEEE.
- [19] Giacomo Cancelli and Mauro Barni, “MPSteg-color: A new steganographic technique for color images,” in *Information Hiding*, Teddy Furon, F. Cayre, G. Doërr, and P. Bas, Eds., Berlin Heidelberg, 2007, vol. 4567 of *Lecture Notes in Computer Science*, pp. 1–15, Springer.
- [20] Matthias Kirchner and Rainer Böhme, “Synthesis of color filter array pattern in digital images,” in *Proceedings of SPIE-IS&T Electronic Imaging: Media Forensics and Security XI*, Edward J. Delp, Jana Dittmann, Nasir D. Memon, and Ping Wah Wong, Eds., San Jose, CA, 2009, vol. 7254, 725421, SPIE.