

Rich Model for Steganalysis of Color Images

Miroslav Goljan and Jessica Fridrich, *Member, IEEE*
Department of ECE
Binghamton University
Binghamton, NY 13902-6000
E-mail: mgoljan, fridrich@binghamton.edu

Rémi Coganne, *Member, IEEE*
ICD - ROSAS - LM2S
Troyes University of Technology
Troyes, France
E-mail: remi.coganne@utt.fr

Abstract—In this paper, we propose an extension of the spatial rich model for steganalysis of color images. The additional features are formed by three-dimensional co-occurrences of residuals computed from all three color channels and their role is to capture dependencies across color channels. These CRMQ1 (color rich model) features are extremely powerful for detection of steganography in images that exhibit traces of color interpolation. Content-adaptive algorithms seem to be hurt much more because of their tendency to modify the same pixels in each channel. The efficiency of the proposed feature set is demonstrated on three different color versions of BOSSbase 1.01 and two steganographic algorithms – the non-adaptive LSB matching and WOW.

Index Terms—Steganalysis, steganography, color, rich models, security

I. INTRODUCTION

Steganography in digital images is nowadays a well advanced discipline. Most studies focus on grayscale images with the unspoken assumption that both steganography design and its detection can be extended in a straightforward manner to color images by considering the color image as three grayscale images. This view, however, ignores the important fact that there exist dependencies among color channels that should be taken into consideration for both steganography and its detection.

It is true that virtually all steganalysis techniques designed for grayscales can be applied to color images, for example, by considering the color image as a three times larger grayscale image or by steganalyzing each channel separately and fusing the detection evidence from each channel. Such methods, however, do not utilize the inter-channel dependencies that can potentially be utilized for detection. The first steganalysis techniques specifically designed for color images were developed for palette images. In [12], the authors pointed out that numerous steganography tools at that time introduced easily identifiable artifacts into the color palette, at least for images with a low color depth (e.g., smaller than 256). A suspiciously large number of close color pairs was proposed in [6] as a distinguishing feature for LSB (Least Significant Bit) replacement and matching in true-color images. The so-called Pairs Analysis [7] (not to be confused with Sample Pairs Analysis) targeted steganographic methods hiding

in palette images by swapping paired up colors (an equivalent of LSB replacement for palette formats). The first steganalysis feature set for color images that considered dependencies among color channels was proposed by Lyu et al. [18]. The authors used higher-order moments of noise residuals obtained using predictors of coefficients in a QMF decomposition of the image from all three color channels. Another detector of LSB matching in color images was proposed in [19] by Westfeld. The author used the 3D color cube representation of the image and proposed the relative frequency of the number of neighbors (in the $3 \times 3 \times 3$ neighborhood) as a steganalysis feature. LSB matching typically introduced a characteristic tail into this histogram, which was not observed in cover images. While this method works reasonably well in decompressed JPEGs, it performs poorly in images with a large number of unique colors, such as never-compressed or resized images. Ker extended the work by Harmsen [9] that used the center of gravity of the histogram characteristic function (HCF) by calibration by resampling and by totalling the color components [13]. Methods based on the HCF, however, do not work well on images that were never compressed. Finally, in a recent work Kirchner et al. [15] improved the Weighted-Stego attack on LSB replacement by using more accurate predictors that consider the dependencies among pixels and channels introduced by the Color Filter Array (CFA) interpolation algorithm.

This work focuses on detection of both non-adaptive LSB matching and modern content-adaptive steganography in true color images in raster formats that were not previously subjected to JPEG compression.¹ In the next section, we describe the common core of all experiments, including the image sources, the classifier used to build detectors, and the tested steganographic schemes. In Section III, we augment an existing feature set designed for grayscale images, the spatial rich model (SRM) computed from the union of all three color channels, with a new set of features formed by co-occurrences between color residuals to capture inter-channel dependencies. Experiments included in Section IV demonstrate that these additional features are extremely efficient in color images with de-

¹Lossy compression suppresses the noise component, which makes steganalysis substantially easier. In decompressed JPEGs, one could also use JPEG compatibility attacks (see, e.g., [16]).

tectable traces of color interpolation and work especially well against content-adaptive algorithms. The experiments are executed for three versions of the color BOSSbase 1.01 database and two steganographic algorithms. Section V contains the summary and a list of additional possible improvements and future directions.

II. COMMON CORE OF EXPERIMENTS

All experiments in this paper are carried out on three versions of BOSSbase 1.01 [1]. Starting with the full-resolution raw images, we converted them using the same script that was used for creating the BOSSbase with the following modifications. The output of `ufraw` (ver. 0.18 with `dcrw` ver. 9.06) was changed to the color ppm format instead of the ppm grayscale. Also, all calls of 'convert' used ppm for the output as well as for resizing so that the smaller image dimension was 512 and for central cropping to 512×512 . As in the original script, the resizing algorithm uses the Lanczos kernel. We thus obtained 10,000 true color 512×512 ppm images. For brevity, we will call this version of color BOSSbase 'BOSSbaseRes' (resized).

The next two versions of BOSSbase were created using the same process as above with two choices for the demosaicking algorithm and with the final resizing operation removed. The crop command was modified to extract the center 512×512 part of the non-resized image. The two demosaicking algorithms used in `dcrw` were the Patterned Pixel Grouping (PPG, sometimes referred to as Pixel Grouping), which is the default algorithm in `dcrw`, and the Adaptive Homogeneity-Directed (AHD) algorithm, which is generally known to provide a better visual quality at the cost of a higher computational complexity. The two corresponding versions of BOSSbase will be called BOSSbasePPG and BOSSbaseAHD.

All detectors were trained as binary classifiers implemented using the FLD ensemble [17] with default settings. A separate classifier was trained for each image source, embedding algorithm, and payload to show how the detection performance depends on the payload size. As described in the original publication, the ensemble by default minimizes the total classification error probability under equal priors $P_E = \min_{P_{FA}} \frac{1}{2}(P_{FA} + P_{MD})$, where P_{FA} and P_{MD} are the false-alarm and missed-detection probabilities. The random subspace dimensionality and the number of base learners is found by minimizing the out-of-bag (OOB) estimate [3] of the testing error, E_{OOB} , on bootstrap samples of the training set. We evaluate the security using the P_E measured on the testing set averaged over ten 5000/5000 database splits denoted as \bar{P}_E . The statistical spread is the sample standard deviation.

Two embedding algorithms were included in this study: the non-adaptive LSB matching (LSBM) with the change rate as the distortion measure and the content-adaptive WOW [10] both simulated at their corresponding rate-distortion bounds. The tested payloads were 0.05, 0.1, 0.2, ..., 0.5 bpc (bits per channel pixel). Both algorithms were

applied to color images by treating their color channels as three grayscale images and embedding the same relative payload in each channel.² A second option that makes sense for the adaptive WOW algorithm is to consider the color image as a three times larger grayscale image and embed the same relative payload in this larger image. This way, WOW could automatically distribute the payload across color channels, potentially embedding more into the noisier channels, thus providing a better security. For investigative purposes and because this paper focuses on detection rather than steganography design, we chose the first option (embedding the same relative payload in each channel) because it allows us to better interpret the results when it comes to statistical detectability of embedding changes across channels.

III. RICH MODEL FOR STEGANALYSIS OF COLOR IMAGES

Most color images that are stored in a raster format, such as TIFF, PNG, BMP, PPM, etc., have undergone a potentially long processing pipeline consisting of gain adjustment (white balance), color correction, color interpolation, gamma correction, denoising, filtering, contrast/brightness adjustment, and profile correction, which includes the correction for lens distortion and chromatic aberration. This is what happens inside the camera that has been set to store images in a true-color uncompressed format (e.g., 24-bit TIFF) and a similar processing chain occurs when a raw image is manually processed in an image editing software, such as Adobe Lightroom. In this case, the resulting image will contain no or only rather weak traces of color interpolation. On the other hand, if a color image exhibits detectable traces of color interpolation, they can likely aid the steganalyst in detection (see the recent work of Kirchner [15] on WS steganalysis). This is why in this paper we test the proposed features on three versions of color BOSSbase (Section II).

Since steganalysis of grayscale images is a well researched topic, we capitalize on existing advancements in feature-based steganalysis. The proposed model for color images is called Spatio-Color rich model or SCRMQ1 and it consists of two different components. The first component is the spatial rich model [8] with a single quantization step $q = 1$ (SRMQ1) with dimensionality 12,753. The SRMQ1 feature is computed for each color channel and the three features are added (merged) to keep the same dimensionality as for grayscale images.³ The second component is a collection of 3D color co-occurrences, the Color Rich Model (CRMQ1), built from the same noise residuals as those used in the SRMQ1 but formed across the three channels of each pixel. We next

²The corresponding change rates for LSB matching were 0.00496, 0.01122, 0.02602, 0.04318, 0.06258, and 0.08398.

³As will become apparent from experiments in Section IV, concatenating the three features instead of adding brings only a negligible improvement.

TABLE I

STEGANALYSIS OF LSB MATCHING AT 0.1 bpc AND WOW AT 0.4 bpc IN BOSSBASERES, BOSSBASEPPG, AND BOSSBASEAHD, THEIR INDIVIDUAL COLOR CHANNELS, AND THEIR GRAYSCALE VERSIONS USING THREE FEATURE SETS: SRMQ1, CRMQ1, AND THEIR UNION SRMQ1.

\overline{P}_E	Dim	LSBM (0.1 bpc)			WOW (0.4 bpc)		
		BOSSbaseRes	BOSSbasePPG	BOSSbaseAHD	BOSSbaseRes	BOSSbasePPG	BOSSbaseAHD
SRMQ1, Merged	12753	.1945 ± .0015	.1490 ± .0016	.1146 ± .0010	.2323 ± .0016	.0225 ± .0004	.0255 ± .0004
SRMQ1, Concat	38259	.1862 ± .0020	.1234 ± .0019	.0936 ± .0009	.2255 ± .0009	.0195 ± .0008	.0157 ± .0008
SRMQ1, Red	12753	.3425 ± .0016	.2756 ± .0014	.2227 ± .0018	.3191 ± .0015	.0885 ± .0014	.0821 ± .0014
SRMQ1, Green	12753	.2097 ± .0015	.1744 ± .0012	.1510 ± .0012	.2487 ± .0016	.0368 ± .0007	.0264 ± .0007
SRMQ1, Blue	12753	.2510 ± .0010	.2579 ± .0018	.2083 ± .0012	.2671 ± .0017	.0811 ± .0009	.0804 ± .0009
SRMQ1, Gray	18157	.1790 ± .0015	.0472 ± .0010	.0297 ± .0008	.1765 ± .0014	.0080 ± .0005	.0041 ± .0005
CRMQ1	5404	.2304 ± .0012	.0621 ± .0007	.0381 ± .0005	.1985 ± .0015	.0117 ± .0004	.0046 ± .0004
SRMQ1, Gray	12753	.2101 ± .0012	.2785 ± .0013	.2373, ±.0011	.2388 ± .0011	.0606 ± .0004	.0440 ± .0005

TABLE II

RELATIVE NUMBER OF PIXELS WITH 2 OR 3 CHANGES IN THEIR COLOR COMPONENTS FOR WOW AT PAYLOADS 0.1, 0.2, AND 0.4 bpc IN BOSSBASERES, BOSSBASEPPG, AND BOSSBASEAHD (ABBREVIATED AS 'RES', 'PPG', AND 'AHD'). THE LAST COLUMN IS FOR LSBM, WHICH IS THE SAME FOR ALL THREE DATABASES.

α	WOW			LSBM
	Res	PPG	AHD	All
0.1	0.0060	0.0023	0.0024	0.0004
0.2	0.0166	0.0072	0.0074	0.0020
0.4	0.0472	0.0246	0.0208	0.0112

describe only those aspects of the SRMQ1 model that are essential for explaining these additional CRMQ1 features while referring the reader to the original publication for more details [8].

The SRMQ1 noise residuals are computed using two types of local pixel predictors – linear, computed by convolving the image with a high-pass filter with a shift-invariant kernel ('spam' residuals), and non-linear ones obtained by taking the maximum / minimum of outputs from several linear filters ('min-max' residuals). The residuals are further divided into five classes depending on the filter structure: filters using first-order, second-order, and third-order differences, edge kernels, and square kernels (see Figure 2 in [8]).

Let us assume that we have a true color image, \mathbf{I} , represented as three $n_1 \times n_2$ matrices $\mathbf{I} = \{\mathbf{R}, \mathbf{G}, \mathbf{B}\}$, $\mathbf{R} = (r_{ij})$, $\mathbf{G} = (g_{ij})$, $\mathbf{B} = (b_{ij})$, $r_{ij}, g_{ij}, b_{ij} \in \{0, 1, \dots, 255\}$, $1 \leq i \leq n_1$, $1 \leq j \leq n_2$. Furthermore, let $\mathbf{Z}^{(r)} = (z_{ij}^{(r)})$ be the quantized noise residual computed only from the red channel, for example, as the average of its horizontal neighbors (denoted as 'spam12h,v' in [8]), $z_{ij}^{(r)} = Q_T(|r_{ij} - (r_{i,j-1} + r_{i,j+1})/2|/q)$, where $Q_T: \mathbb{R} \rightarrow \{-T, -T+1, \dots, T\}$ is a quantizer with $2T+1$ integer centroids and q is the quantization step (here, $q=1$). Similarly, we also compute the residuals, $z_{ij}^{(g)}$, $z_{ij}^{(b)}$ in the other two channels. The SRMQ1 uses 4D co-occurrences of four neighboring residuals, e.g., $(z_{ij}^{(r)}, z_{i,j+1}^{(r)}, z_{i,j+2}^{(r)}, z_{i,j+3}^{(r)})$ for the horizontal co-occurrence. Because the co-occurrences formed in the horizontal and vertical directions capture different statistics among pixels, they are treated sepa-

rately in SRMQ1. In the CRMQ1 model, however, we form the co-occurrences across the color channels for each pixel from the triplets $(z_{ij}^{(r)}, z_{ij}^{(g)}, z_{ij}^{(b)})$. Thus, we only need one co-occurrence even when the noise residual is "directional" as is the case in the example above. Due to symmetries of natural images, we combine the 3D co-occurrences computed from the horizontal version of the predictor with the vertical version: $z_{ij}^{(r)} = Q_T(|r_{ij} - (r_{i-1,j} + r_{i+1,j})/2|/q)$. Depending on the symmetries of the residual kernel, we may be adding statistics from four or even eight residual versions after including the mirror versions of the kernels and their versions rotated by 90 degrees (e.g., see the residuals 'minmax24' and 'minmax48h,v' in [8]).

Since the SRMQ1 is formed from 8 first-order and 8 third-order residuals, 5 second-order residuals, 8 edge residuals, and two square residuals, there will be a total of 31 3D co-occurrences in the CRMQ1 model. They all will be symmetrized by direction and by sign as in the original SRM (sign symmetry means merging the bin d_1, d_2, d_3, d_4 with $-d_1, -d_2, -d_3, -d_4$, $d_k \in \{-T, \dots, T\}$, while a directional symmetry merges the bin with its mirror version d_4, d_3, d_2, d_1). Because we form 3D co-occurrences in CRMQ1, we opted for a larger threshold $T=3$ instead of $T=2$ used for the residuals in the SRM. After the symmetrization, all 3D co-occurrences have the dimensionality of 100 for spam type residuals and 196 for min-max residuals. Because there are 7 spam type residuals and 24 min-max residuals, the final dimensionality of the CRMQ1 model is thus $7 \times 100 + 24 \times 196 = 5404$.

To summarize, the proposed feature set consists of the 12,753 dimensional SRMQ1 model (that uses 4D co-occurrences with $T=2$ and quantization $q=1$) merged over the color channels and the 5,404 dimensional CRMQ1 model formed by 3D co-occurrences across color channels with $T=3$ and $q=1$. We call this feature set the SRMQ1 model ('C' for the Color version of the SRMQ1) and its dimensionality is $12753 + 5404 = 18157$.

IV. EXPERIMENTS

The results of our first experiment are displayed in Table I, which shows the average detection error \overline{P}_E for one payload, and in Figures 2 and 3 showing \overline{P}_E as a

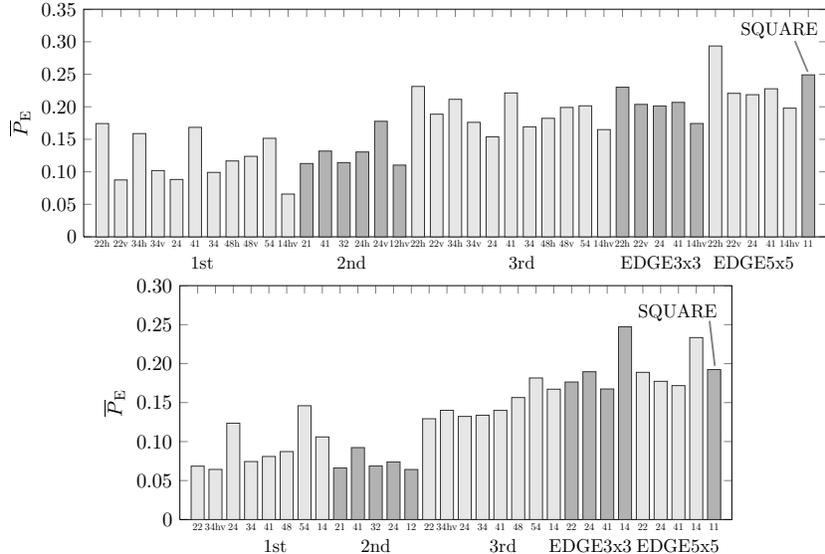


Fig. 1. Average detection error of individual submodels in the SRMQ1 (top) and CRMQ1 (bottom) for WOW at 0.4 bpc in BOSSbasePPG. The submodels are labeled as in [8]. The detection errors of SRMQ1, CRMQ1, and SCRMQ1 are 0.0225, 0.0117, and 0.0080, respectively. Note that submodels with smaller kernels generally detect better than large kernels.

function of payload for LSBM and WOW in all three databases. In particular, we wish to point out the following quite interesting insights.

The table shows that concatenating the SRMQ1 features as opposed to merging them (and saving a factor of 3 on feature dimensionality) brings only a small improvement ($\approx 2.5\%$). Also, as expected, steganalysis using SCRMQ1 is easier in images with traces of demosaicking than in resized images for both LSBM and WOW.

While for LSB matching in resized images the CRMQ1 model improves the detection only marginally, it is more effective for the content-adaptive WOW (observe the difference between SRMQ1 and SCRMQ1 in top charts of Figures 2 and 3). We explain this by noting that the adaptive WOW is much more likely to modify more than one channel in a given pixel than the non-adaptive LSBM (see Table II), which is better captured by the CRMQ1 model. For both LSBM and WOW, the CRMQ1 features are now much stronger than the SRMQ1 and are responsible for most of the detection power in SCRMQ1 apparently capitalizing on demosaicking artifacts (see also Figure 1). Moreover, WOW appears to be hurt by the power of the CRMQ1 model significantly more than LSBM for the same reason as stated above. This high detectability of WOW is, of course, due to the naive way of applying WOW to color images. The security of content-adaptive algorithms in color images can likely be significantly improved by properly designing the distortion function to consider inter-channel dependencies.

The choice of the demosaicking algorithm between PPG and AHD does not seem to play a significant role for steganalysis using either feature set. The images demosaicked using the AHD algorithm appear slightly eas-

ier to steganalyze than the images obtained using the PPG. Surprisingly, when converting the color images in BOSSbaseRes to grayscale using Matlab's 'rgb2gray' the detectability using the spatial SRMQ1 model appears almost identical (the top graphs in Figures 2 and 3). Apparently, the decrease in detectability due to the effect of the square root law [5] when steganalyzing in three times smaller images is almost exactly compensated by the smoothing character of converting the three color channels to grayscale. Of course, in images with traces of demosaicking the difference between steganalyzing with SRMQ1 in color and in the corresponding grayscale images is much bigger because the conversion to grayscale largely eliminates the color interpolation artifacts.

The detection in the individual color channels is the easiest in the green and hardest in red. This is consistently true for all image sources and both algorithms. This is probably because the red channel is always the noisiest channel. The root mean square (RMS) of the noise residual obtained using the KB predictor [2] in the R, G, and B color channels averaged across the entire database was [26.5, 24.9, 24.8] in BOSSbaseRes and [16.9, 12.1, 16.5] in BOSSbasePPG, confirming that the red channel is always the noisiest. We attribute this increased noise level in the red channel to a higher gain adjustment (white balance coefficient) for the red channel. Since the white balance gain used for each channel can be extracted from the EXIF header, we confirmed this by computing the average ratio of the gains $g_{red}/g_{green} \approx 2.03$ and $g_{blue}/g_{green} \approx 1.47$ over the entire database.

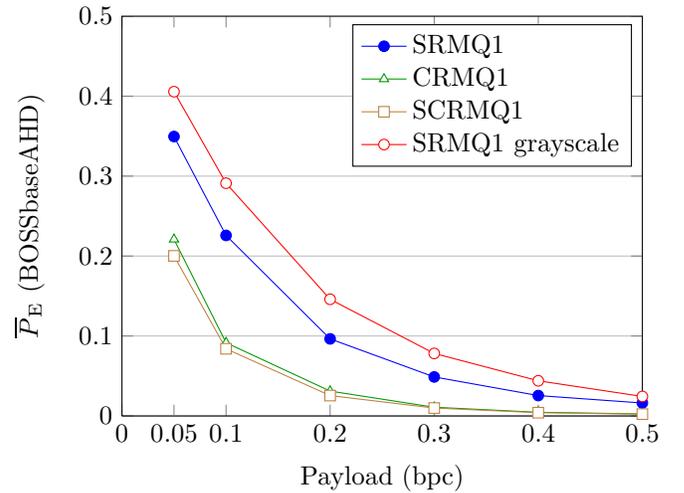
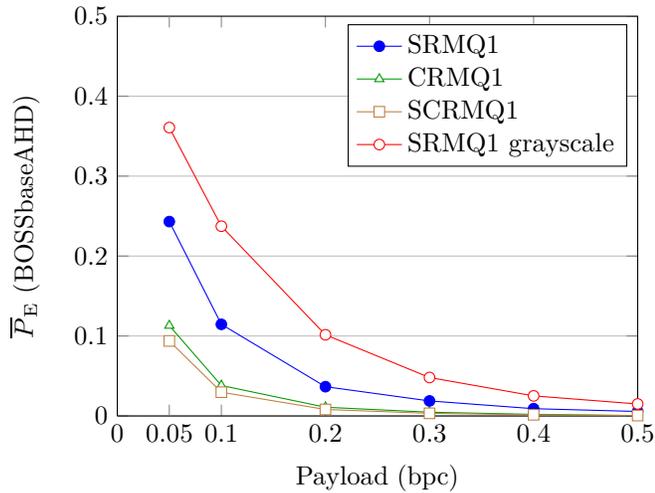
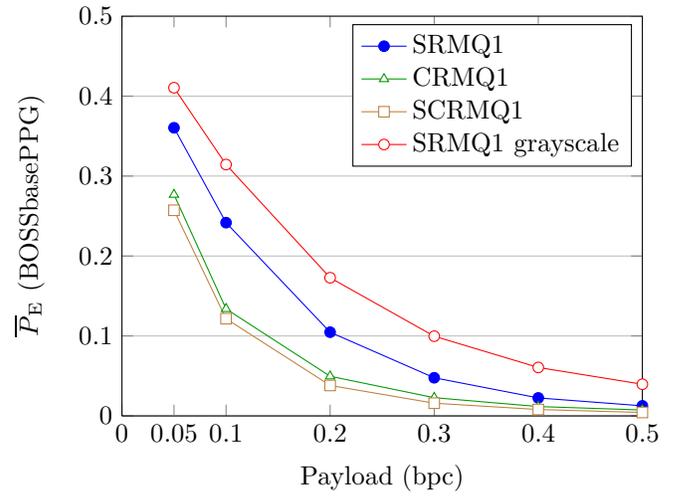
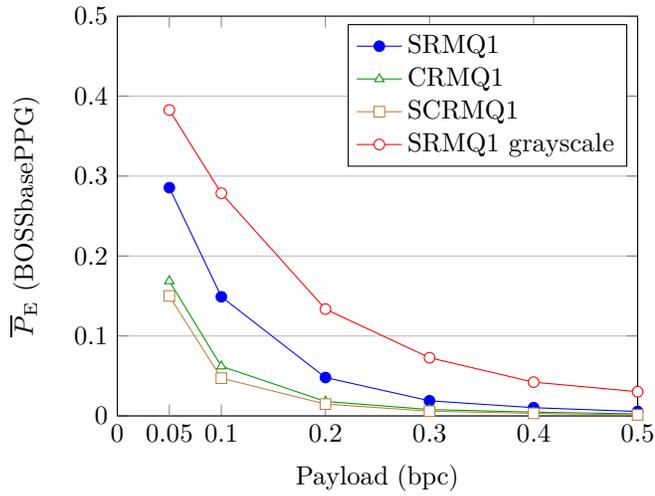
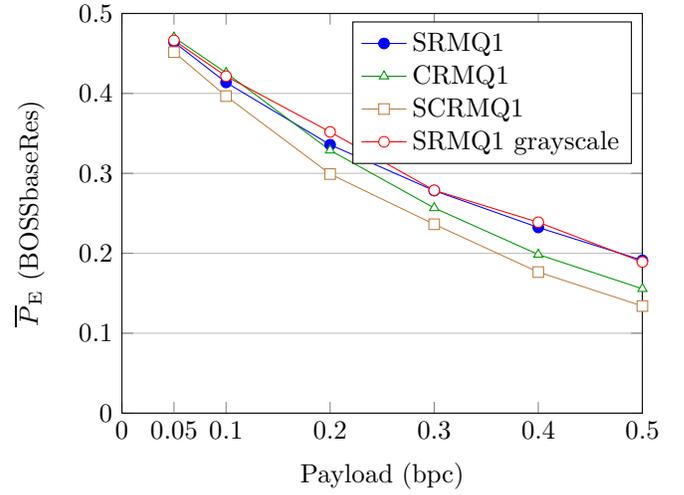
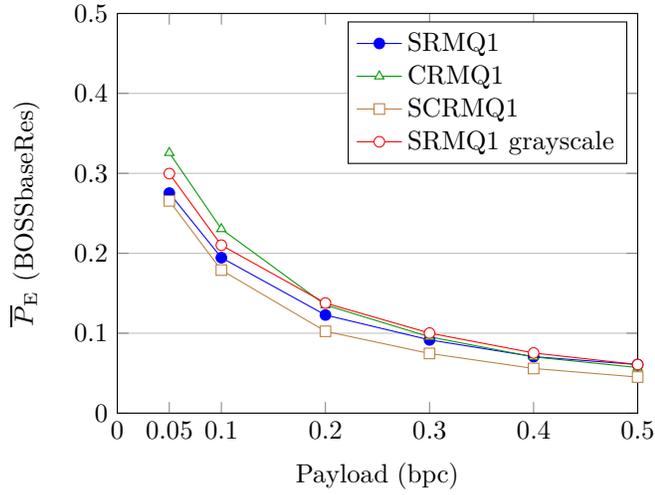


Fig. 2. Detection error \bar{P}_E as a function of payload for LSBM. From top to bottom: BOSSbaseRes, BOSSbasePPG, and BOSSbaseAHD.

Fig. 3. Detection error \bar{P}_E as a function of payload for WOW. From top to bottom: BOSSbaseRes, BOSSbasePPG, and BOSSbaseAHD.

While in resized images the best submodels are those with large supports, in images with traces of demosaicking, the tendency is the opposite (note the detection of WOW at 0.4 bpc in BOSSbasePPG using the individual submodels of SCRMQ1 in Figure 1). This is most likely

due to the fact that each color channel is an interlaced union of four types of pixels based on their position w.r.t. the Bayer CFA. Residuals with a larger support mix these heterogeneous signals more than residuals with a smaller support, resulting in a weaker detection.

V. CONCLUSION

This paper proposes an extension of the spatial rich model (SRM), which was designed for grayscale images, to allow for more accurate detection of steganography in color images. The SRMQ1 features are augmented by a collection of symmetrized 3D co-occurrences of residuals between color channels – the Color Rich Model (CRMQ1) of dimension 5404. While these additional features help detection only marginally in color images that do not exhibit traces of color interpolation, when these traces are present the CRMQ1 features become responsible for the majority of the detection power independently of the two tested demosaicking algorithms and the steganographic algorithms. Curiously, in demosaicked images the difference between the security of the LSB matching and the content-adaptive WOW when naively applied independently to each color channel becomes rather small due to the fact that WOW concentrates the embedding changes in textured regions and thus more often modifies two or more color channels of any given pixel. Content-adaptive steganography for color images obviously needs to use distortion functions that can capture dependencies across the color channels.

We note that further obvious improvements in detection are possible by investigating the best choice of the quantization step in SRM with a single quantization step, by utilizing the full version of SRM with all three quantization steps, or by employing the projection SRM (PSRM) [11]. Additionally, the rich model could be compactified to a lower dimension and more powerful classifiers, such as the Gaussian SVM can be utilized to further boost the detection. The SRM part of the model could be compactified by feature selection or by employing variable co-occurrences [4]. We observed that the CRMQ1 features' dimensionality can be decreased from 5404 to 2074 by using a smaller threshold $T = 2$ with only a small performance penalty. We do not explore these directions in this paper due to the limited space and because we prefer to focus on introducing qualitatively new types of features and on revealing interesting phenomena in steganalysis of color images rather than on the best achievable performance.

In images with demosaicking artifacts, steganography detection greatly benefits from the proposed co-occurrences between the residuals of individual color channels. The residuals, though, are still computed from each color channel from all pixels and are thus “blind to the CFA.” In our future effort, we plan to investigate the idea of making the residuals themselves aware of the underlying structure of the CFA to better separate the signals with different statistical properties and thus further improve detection. Since color interpolation artifacts in digital images can be rather reliably detected [14], the steganalyst may use a preclassifier when testing a given image and send it to the appropriately trained classifier.

The code for SRMQ1 and CRMQ1 is available from <http://dde.binghamton.edu/download/>.

VI. ACKNOWLEDGEMENTS

This work was supported by Air Force Office of Scientific Research under the research grant number FA9950-12-1-0124. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation there on. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied of AFOSR or the U.S. Government. Rémi Coganne was also partially funded by Troyes University of Technology (UTT) strategic program COLUMBO and STEG-DETECT program for scholar mobility. This work was done by Rémi when he was a visiting scholar at Binghamton University.

REFERENCES

- [1] P. Bas, T. Filler, and T. Pevný. Break our steganographic system – the ins and outs of organizing BOSS. In *Proc. IH Conf.*, volume 6958 of *LNCS*, pages 59–70, 2011.
- [2] R. Böhme. Weighted stego-image steganalysis for JPEG covers. In *Proc. IH Conf.*, volume 5284 of *LNCS*, pages 178–194, 2007.
- [3] L. Breiman. Bagging predictors. *Machine Learning*, 24:123–140, August 1996.
- [4] L. Chen, Y. Q. Shi, P. Sutthiwan, and X. Niu. A novel mapping scheme for steganalysis. In *Proc. IWDW*, volume 7809 of *LNCS*, pages 19–33, 2012.
- [5] T. Filler, A. D. Ker, and J. Fridrich. The Square Root Law of steganographic capacity for Markov covers. In *Proc. SPIE, EI*, vol. 7254, pages 08 1–11, 2009.
- [6] J. Fridrich, R. Du, and M. Long. Steganalysis of LSB encoding in color images. volume 3, pages 1279–1282, 2000.
- [7] J. Fridrich, M. Goljan, and D. Soukal. Higher-order statistical steganalysis of palette images. In *Proc. SPIE, EI*, vol. 5020, pages 178–190, 2003.
- [8] J. Fridrich and J. Kodovský. Rich models for steganalysis of digital images. *Trans. IEEE TIFS*, 7(3):868–882, 2011.
- [9] J. J. Harmsen and W. A. Pearlman. Steganalysis of additive noise modelable information hiding. In *Proc. SPIE, EI*, vol. 5020, pages 131–142, 2003.
- [10] V. Holub and J. Fridrich. Designing steganographic distortion using directional filters. In *Proc. IEEE WIFS*, 2012.
- [11] V. Holub and J. Fridrich. Random projections of residuals for digital image steganalysis. *Trans. IEEE TIFS*, 8(12):1996–2006, 2013.
- [12] N. F. Johnson and S. Jajodia. Steganalysis of images created using current steganography software. In *Proc. IHW*, volume 1525 of *LNCS*, pages 273–289, 1998.
- [13] A. D. Ker. Resampling and the detection of LSB matching in color bitmaps. In *Proc. SPIE, EI*, vol. 5681, pages 1–15, 2005.
- [14] M. Kirchner. Efficient estimation of CFA pattern configuration in digital camera images. In *Proc. SPIE, EI*, vol. 7541, pages 11 1–12, 2010.
- [15] M. Kirchner and R. Böhme. Steganalysis in technicolor: Boosting WS detection of stego images from CFA-interpolated covers. In *Proc. IEEE ICASSP*, 2014.
- [16] J. Kodovský and J. Fridrich. JPEG-compatibility steganalysis using block-histogram of recompression artifacts. In *Proc. IH Conf.*, volume 7692 of *LNCS*, pages 78–93, 2012.
- [17] J. Kodovský, J. Fridrich, and V. Holub. Ensemble classifiers for steganalysis of digital media. *Trans. IEEE TIFS*, 7(2):432–444, 2012.
- [18] S. Lyu and H. Farid. Steganalysis using color wavelet statistics and one-class support vector machines. In *Proc. SPIE, EI*, vol. 5306, pages 35–45, 2004.
- [19] A. Westfeld. Detecting low embedding rates. In *Proc. IHW*, volume 2578 of *LNCS*, pages 324–339, 2002.