

Side-Informed Steganography with Additive Distortion

Tomáš Denemark and Jessica Fridrich, *Senior Member, IEEE*

Department of ECE
Binghamton University
Binghamton, NY 13902-6000
{tdenema1,fridrich}@binghamton.edu

Abstract—Side-informed steganography is a term used for embedding secret messages while utilizing a higher quality form of the cover object called the precover. The embedding algorithm typically makes use of the quantization errors available when converting the precover to a lower quality cover object. Virtually all previously proposed side-informed steganographic schemes were limited to the case when the side-information is in the form of an uncompressed image and the embedding uses the unquantized DCT coefficients to improve the security when JPEG compressing the precover. Inspired by the side-informed (SI) UNIWARD embedding scheme, in this paper we describe a general principle for incorporating the side-information in any steganographic scheme designed to minimize embedding distortion. Further improvement in security is obtained by allowing a ternary embedding operation instead of binary and computing the costs from the unquantized cover. The usefulness of the proposed embedding paradigm is demonstrated on a wide spectrum of various information-reducing image processing operations, including image downsampling, color depth reduction, and filtering. Side-information appears to improve empirical security of existing embedding schemes by a rather large margin.

I. INTRODUCTION

The goal of steganography is to communicate secret messages to another party by hiding the secrets in cover objects so that the Warden, who monitors the traffic, cannot distinguish between genuine cover objects and objects carrying secret data. In steganography by cover modification, the secret is embedded by making changes to the cover. If the cover-source distribution is known and available to the communicating parties as well as the Warden, the rate of perfectly secure steganographic communication is positive [24] even when the actions of both the sender and the (possibly active) Warden are power limited. When the cover source is empirical (incognizable) in nature, such as digital media, the individual cover elements exhibit complex dependencies that are highly non stationary. This means that in practice neither the steganographer or the Warden know the model and have to work with approximations. This has fundamental consequences for the steganographer, who is now unable to achieve perfect security. With increasing number of communicated information, the sender has to embed with

a vanishing rate to curb the risk of the Warden detecting the usage of steganography [16], [5].

To alleviate the lack of the cover model, some steganographic schemes make use of the knowledge of the so-called precover.¹ The precover is usually a higher-quality representation of the cover, such as the raw image before it is converted to a true-color image, such as TIFF or JPEG. Historically, the first side-informed embedding scheme was the embedding-while-dithering steganography [6], in which the secret message was hidden in selected pixels of a GIF image by perturbing the quantization to palette colors and dithering both the quantization error and the embedding distortion. Another early example is the Perturbed Quantization [7], which hides secrets during a recompression of a JPEG file. Side-informed schemes gained on popularity with the introduction of MMEx [17] and BCHopt [22], both designed to hide messages in JPEG files while utilizing the rounding errors of DCT coefficients. Such schemes and their improvements [23], [14], [10] offered a significant increase in empirical security when compared with embedding schemes that do not use any side-information [18]. Finally, in [13] the authors described a side-informed JPEG steganographic scheme called SI-UNIWARD, which is currently among the most secure algorithms available for JPEG images [12].

In the next section, we propose a rather simple way how to incorporate quantization errors obtained when processing the precover to its cover form within any steganographic scheme designed to minimize an additive embedding distortion. Then, we further generalize the approach to make use of the more powerful ternary embedding operation instead of a binary one, which is what all side-informed schemes have traditionally used. The effectiveness of the proposed methodology is demonstrated in Section III, where we carry out experiments with several information-reducing operations in both the spatial and JPEG domain. The paper is concluded in Section IV.

II. INCORPORATING SIDE-INFORMATION

In this section, we introduce a simple idea how to incorporate side-information in any steganographic scheme that

¹The concept of precover is due to Ker [15].

minimizes additive distortion. We specifically discuss two novel aspects, which include the departure from a binary embedding operation to ternary and the computation of costs from the unquantized cover.

The following notational conventions are adopted in this paper. The index pair ij will always be used for pixels or DCT coefficients and uv for wavelet coefficients. Boldface font is used for matrices and vectors.

We will recognize three types of images – a precover image \mathbf{P} , unquantized cover \mathbf{U} , and cover \mathbf{X} . The cover is obtained from the precover using some information-reducing operation that involves quantization as the last step. Even though the proposed approach can certainly be applied to color cover images, for simplicity of the exposition, we will assume that $\mathbf{X} = (X_{ij}) \in \mathcal{I}_L^{n_1 \times n_2}$, $\mathcal{I}_L = \{0, 1, \dots, 2^L - 1\}$, is an L -bit grayscale image with $n_1 \times n_2$ pixels. At this point, we refrain from formalizing the concept of the precover and merely state that it is some higher quality version of the cover. The precover may be in a different format than the cover, it may have a higher color depth, and may be larger than $n_1 \times n_2$ pixels. In this paper, we also allow the precover to be color. We assume that there be a transformation T that maps the precover \mathbf{P} to $\mathbb{R}^{n_1 \times n_2}$ such that $\mathbf{X} = Q_L(T(\mathbf{P}))$ with $\mathbf{U} = T(\mathbf{P})$ the unquantized cover, where Q_L is a quantizer with 2^L centroids \mathcal{I}_L . Symbolically,

$$\mathbf{P} \xrightarrow{T} \mathbf{U} \xrightarrow{Q_L} \mathbf{X}. \quad (1)$$

Furthermore, we will assume that we have a steganographic scheme \mathcal{A} designed to embed while minimizing an additive distortion function. This is currently the most successful paradigm for constructing steganographic schemes in any domain, including those based on non-additive distortion [3], [2], [21]. Such steganography typically starts with computing the costs ($\rho_{ij}^{(A)}$) of changing each cover element X_{ij} . In this article, we will use only additive schemes \mathcal{A} where the cost of changing the X_{ij} by 1 and by -1 are the same. Again, most additive embedding schemes do possess this property, e.g., S-UNIWARD, J-UNIWARD [13], HILL [20], WOW [11], and UED [9]. The cost $\rho_{ij}^{(A)}$ typically depends on some local neighborhood of pixel X_{ij} . At this point, we note that in all schemes known to the authors it is possible to compute the costs from the unquantized cover \mathbf{U} instead of the cover \mathbf{X} . To distinguish such costs, we will explicitly mark this dependency: $\rho_{ij}^{(A)}(\mathbf{U})$ or $\rho_{ij}^{(A)}(\mathbf{X})$.

The side-information that our proposed general embedding scheme will utilize is the unquantized cover \mathbf{U} . The quantization error due to applying the quantizer Q_L will be denoted

$$e_{ij} = U_{ij} - X_{ij} = U_{ij} - Q_L(U_{ij}). \quad (2)$$

If the cover element is modified from X_{ij} to Y_{ij} , the total distortion due to quantization and embedding with

respect to the unquantized cover is thus

$$e'_{ij} = U_{ij} - Y_{ij}. \quad (3)$$

The costs $\rho_{ij}^{(\text{SI})}$ of the side-informed version of the embedding algorithm \mathcal{A} are obtained by modulating the original costs by the difference $|e'_{ij}| - |e_{ij}|$:

$$\rho_{ij}^{(\text{SI})} = (|e'_{ij}| - |e_{ij}|)\rho_{ij}^{(A)}. \quad (4)$$

Note that $\rho_{ij}^{(\text{SI})} \geq 0$ because $|e_{ij}| \leq |e'_{ij}|$ for all ij as e_{ij} is the smallest amount U_{ij} can be modified to obtain a plausible cover value (a value from \mathcal{I}_L). The modulation in (4) makes intuitive sense because the new costs reflect not only the local image complexity but also take into account the distortion w.r.t. the unquantized cover. The hope is that by minimizing this distortion, the embedding will disturb the statistical properties of covers less.

At this point, we make a comparison to previous art. Virtually all previously proposed side-informed schemes were restricted to binary embedding operations. During embedding, the value U_{ij} was either quantized to $Y_{ij} = X_{ij} = Q_L(U_{ij})$ or it was rounded “to the other side” $Y_{ij} = X_{ij} + \text{sign}(U_{ij} - X_{ij}) = U_{ij} + \text{sign}(e_{ij}) - e_{ij}$, which means the embedding operation was inherently binary. In this case, $|e'_{ij}| - |e_{ij}| = 1 - 2|e_{ij}|$.

In this article, we allow ternary side-informed embedding. Indeed, when $e_{ij} \approx 0$, there is no reason to restrict the embedding to a binary operation as rounding to $Y_{ij} = X_{ij} + \text{sign}(e_{ij})$ becomes almost as expensive ($|e'_{ij}| - |e_{ij}| = 1 - 2|e_{ij}|$) as rounding to $Y_{ij} = X_{ij} - \text{sign}(e_{ij})$ ($|e'_{ij}| - |e_{ij}| = 1 + |e_{ij}| - |e_{ij}| = 1$). Thus, when using ternary embedding in our side-informed steganography, the costs of changing X_{ij} by ± 1 are not equal:

$$\rho_{ij}^{(\text{SI})+} = (1 - 2|e_{ij}|)\rho_{ij}^{(A)} \text{ if } Y_{ij} = X_{ij} + \text{sign}(e_{ij}), \quad (5)$$

$$\rho_{ij}^{(\text{SI})-} = \rho_{ij}^{(A)} \text{ if } Y_{ij} = X_{ij} - \text{sign}(e_{ij}). \quad (6)$$

The actual embedding needs to be executed with the multi-layered version of syndrome-trellis codes (STCs) [4]. An embedding simulator will change pixel X_{ij} by $\pm \text{sign}(e_{ij})$ with probabilities

$$\beta_{ij}^{(\pm)} = \frac{e^{-\lambda \rho_{ij}^{(\text{SI})\pm}}}{1 + e^{-\lambda \rho_{ij}^{(\text{SI})+}} + e^{-\lambda \rho_{ij}^{(\text{SI})-}}}, \quad (7)$$

with $\lambda > 0$ determined by the payload length M (in bits):

$$M = \sum_{i,j} -\beta_{ij}^+ \log_2 \beta_{ij}^+ - \beta_{ij}^- \log_2 \beta_{ij}^- - (1 - \beta_{ij}^+ - \beta_{ij}^-) \log_2 (1 - \beta_{ij}^+ - \beta_{ij}^-). \quad (8)$$

A. Discussion and relationship to prior art

The modulation of costs by the difference $|e'_{ij}| - |e_{ij}|$ has been proposed in the past. The BCHopt embedding scheme [22] for JPEG images modulates the costs in the form of the quantization steps with this factor. The same factor also appears in EBS [23], NPQ [14],

UED [10], and SI-UNIWARD [13]. Here, we note that the description of SI-UNIWARD as appeared in [13] does not correspond to the actual implementation available on the authors’ web site (http://dde.binghamton.edu/download/stego_algorithms/). We elaborate on this issue in the appendix. What is new in our proposal is using the factor to modulate the costs of any additive steganography, for example, in the spatial domain. Additionally, we propose two more innovations – the ternary embedding operation and we also compute the costs of \mathcal{A} from the unquantized cover \mathbf{U} rather than the cover \mathbf{X} . It is shown in the next section that both further improve the empirical security.

III. EXPERIMENTS

The precover source for all our experiments was the BOSSbase 1.01 [1] database with images in their full resolution RAW format. We used a script that utilized ‘ufraw’ to convert them to the RGB TIFF format of the same resolution. This included gain adjustment, gamma correction, and color interpolation. All subsequent processing was done in Matlab rather than ImageMagick to obtain an easy access to the non-rounded values. The final quantizer Q_L used $L = 8$ for spatial domain and $L = 11$ for experiments in the JPEG domain. For brevity, the transformation T will be described symbolically by arrows between different image representations in the form $(\text{color})_{\text{type}}^{\text{size}}$, where $\text{color} \in \{\text{RAW}, \text{RGB}, \text{GRAY}\}$, $\text{size} \in \{\text{FULL}, 512^2\}$ for full-size and 512×512 images, and $\text{type} \in \{n\text{B}, \text{DBL}\}$ for n bit integers and doubles.

A side-informed scheme based on embedding algorithm \mathcal{A} that uses q -ary embedding operation and computes the costs from image $\mathbf{C} \in \{\mathbf{U}, \mathbf{X}\}$ will be denoted as SI q - \mathcal{A} - \mathbf{C} .

All detectors were trained as binary classifiers implemented using the FLD ensemble [19] with default settings. The security is evaluated using the ensemble’s ‘out-of-bag’ (OOB) error E_{OOB} averaged over ten ensemble runs with different seeds. In the spatial domain, we steganalyze with the SRM features [8] while PHARM features [12] were used for the JPEG domain.

A. Spatial domain

In this section, we investigate the empirical security of side-informed HILL [20] and S-UNIWARD [13] when \mathbf{U} represents non-rounded pixel values. The goal of the three experiments below is to investigate the effect of the transformation T , the difference between binary and ternary embedding, and between computing the costs from \mathbf{U} and \mathbf{X} . Because our cover images were obtained using operations from Matlab instead of ImageMagick, our cover source differs from the original BOSSbase 1.01, and the detection errors will be different than the ones reported in [20] and [13]. This is not an issue as we are primarily interested in the relative improvement of the side-informed schemes over the original algorithms.

Table I
MEAN E_{OOB} FOR HILL, S-UNIWARD AND THEIR SIDE-INFORMED VARIANTS WHEN UTILIZING THE QUANTIZATION ERROR AFTER RESIZING WITH DIFFERENT KERNELS AT 0.4 BPP.

	bilinear	bicubic	box	triangle	cubic	Lanczos2	Lanczos3
HILL	0.0978	0.1564	0.2038	0.0981	0.1543	0.1539	0.1684
SI3-HILL-U	0.1731	0.2411	0.2899	0.1738	0.2421	0.2462	0.2652
SUNI	0.0646	0.1092	0.1249	0.0651	0.1072	0.1081	0.1233
SI3-SUNI-U	0.1261	0.1941	0.2562	0.1262	0.1916	0.1935	0.2061

1) Resizing:

$$T : (\text{RAW})^{\text{FULL}} \longrightarrow (\text{RGB})_{\text{8B}}^{\text{FULL}} \xrightarrow{\text{gray}} (\text{GRAY})_{\text{8B}}^{\text{FULL}} \xrightarrow[\text{crop}]{\text{resize}} (\text{GRAY})_{\text{DBL}}^{512^2}$$

The results for HILL and S-UNIWARD (Fig. 1) point out two important facts. First, the choice between computing the costs from \mathbf{U} and \mathbf{X} has a small effect on the overall detection because the costs of both algorithms are insensitive to small perturbations of the cover. Thus, in all our subsequent experiments we use just the costs computed from the unquantized cover, $\rho_{ij}(\mathbf{U})$. Second, the gain in security when using the ternary embedding over the binary is significant. Third, the side-informed schemes achieve significantly higher security despite making more embedding changes (Fig. 2) and embedding into smooth areas (Fig. 3). This means that the security of SI schemes solely hinges upon the difficulty of estimating the rounding errors from the quantized (and embedded) image. Finally, Table I shows the effect of the resizing kernel in Matlab’s ‘imresize’ for SI-HILL and SI-S-UNIWARD at 0.4 bpp. The filter ‘nearest’ is missing as it does not produce any rounding error.

2) Color conversion:

$$T : (\text{RAW})^{\text{FULL}} \longrightarrow (\text{RGB})_{\text{8B}}^{\text{FULL}} \xrightarrow[\text{crop}]{\text{resize}} (\text{RGB})_{\text{8B}}^{512^2} \xrightarrow{\text{gray}} (\text{GRAY})_{\text{DBL}}^{512^2}$$

In Fig. 4, we compare HILL and S-UNIWARD and their side-informed variants when utilizing the color conversion rounding error. The images are first resized so that the smaller side is 512 pixels and then cropped to square. The conversion used a linear combination of individual RGB channels with coefficients [0.2989, 0.5870, 0.1140]. This experiment again shows a strong positive influence of the ternary embedding operation, increasing E_{OOB} by almost 10%.

3) Quantization:

$$T : (\text{RAW})^{\text{FULL}} \longrightarrow (\text{RGB})_{\text{16B}}^{\text{FULL}} \xrightarrow{\text{crop}} (\text{RGB})_{\text{16B}}^{512^2} \xrightarrow{\text{gray}} (\text{GRAY})_{\text{DBL}}^{512^2}$$

In this case, the transformation T does not include resizing, which makes the cover images smoother than in the previous experiments. Thus, the steganographic algorithms adapt to the acquisition noise naturally present in the image rather than the content. It is rather interesting that for this source HILL and S-UNIWARD have almost identical performance (Fig. 5). The gain of ternary embedding is mostly pronounced for medium payloads.

B. JPEG domain

The precover source for all experiments in this section was the BOSSbase 1.01 database of 512×512 grayscale images in the PGM format, which were JPEG compressed with Matlab’s ‘imwrite’ with 75% quality factor. Here,




Figure 1. Mean E_{OOB} for HILL (top) and S-UNIWARD (bottom) and their SI versions with the quantization error after resizing with Lanczos 3 kernel as the side-information when computing the costs from the unquantized and quantized cover.

the unquantized cover \mathbf{U} are the non-rounded DCT coefficients divided by the corresponding quantization steps.

Fig. 6 shows the security of J-UNIWARD and its side-informed variants. Note that the SI embedding in JPEGs exhibits very different properties than in the spatial domain. The difference between the costs $\rho_{ij}(\mathbf{U})$ and $\rho_{ij}(\mathbf{X})$ is now more influential while the choice of the embedding operation (binary vs. ternary) is small with the ternary embedding giving a slightly worse performance. Both observations can be attributed to the much larger quantization step. Indeed, the harsher quantization removes more information about the cover source, which makes the costs computed from the unquantized cover more tightly related to detectability. On the other hand, the larger quantization step makes the cost of ternary embedding also higher than in the spatial domain. Finally, note that the actual implementation of SI-UNIWARD [13] corresponds to SI2-JUNI-U (Appendix A).

IV. CONCLUSIONS

Side-informed steganography has been studied in the past but was limited only to JPEG and palette images. In this paper, we formalize a general principle for utilizing side-information in any steganographic scheme that minimizes distortion, further enhance security by allowing




Figure 2. By rows: cover image, its detail, embedding changes for HILL and SI3-HILL-U at 0.4 bpp for resizing with Lanczos 3 kernel.




Figure 3. Change rate for HILL and S-UNIWARD and their SI versions when resizing with Lanczos 3 kernel. The values are averages over all 10,000 images in our source.

ternary embedding, and investigate the impact of computing embedding costs from quantized and unquantized covers. The investigation is experimental and carried out for resizing, color depth reduction, and color to grayscale conversion in the spatial domain and for quantization during JPEG compression. The gain in empirical security and the effect of the proposed measures appears to depend mainly on the ratio of the quantization step used for the final quantization and the image dynamic range. In the spatial domain, this ratio is small, which makes the effect of computing the costs from the unquantized cover rather than the quantized cover negligible. This is also because the selection channel of modern spatial domain embedding schemes is insensitive to small perturbations. On the other




Figure 4. Mean E_{OOB} for HILL, S-UNIWARD and their SI versions with the quantization error after RGB to grayscale conversion as the side-information.

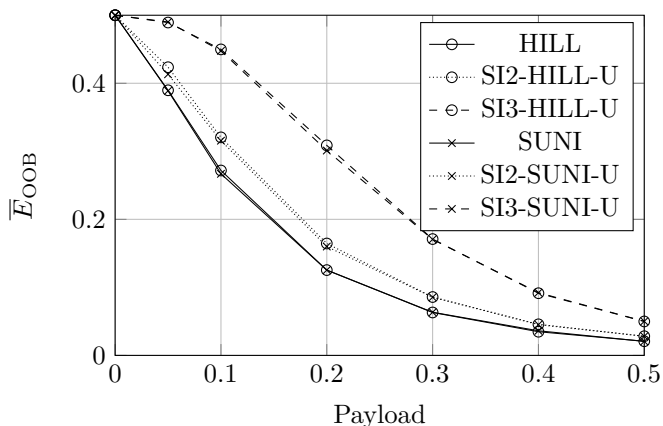


Figure 5. Mean E_{OOB} for SRM for HILL, S-UNIWARD and their SI versions with the quantization error after color depth reduction as the side-information.

hand, the effect of allowing a ternary embedding operation is quite significant because rounding “to both sides” is less expensive due to the fine quantization step. The situation is exactly the opposite for JPEG images. There, the ternary embedding does not bring any improvement due to the large amplitude of embedding changes while the costs computed from the unquantized precover give better security as more information about the precover source is lost due to the much harsher quantization.

The next obvious question is the security of SI schemes against selection-channel-aware steganalysis. To facilitate this, we need to develop techniques for estimating the rounding errors from the quantized and embedded image and design selection-channel-aware feature sets for steganalysis in the JPEG domain. We hypothesize that even inaccurate estimation of the rounding errors in combination with features capable of incorporating this information may decrease the security gain of SI schemes depending on the estimation accuracy. In particular, it may be feasible to obtain reasonable estimates of the quantization

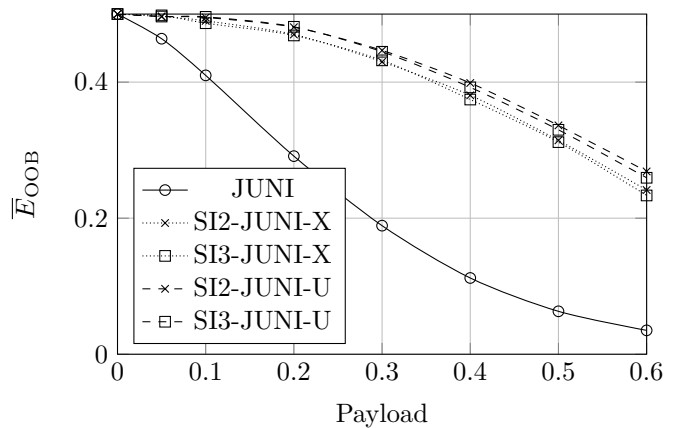


Figure 6. Mean E_{OOB} for J-UNIWARD and its side-informed variants when utilizing the quantization error after JPEG compression when computing the costs from the unquantized cover and the cover. The used quality factor is 75%.

errors in modelable segments, such as portions of blue sky, when the quantization step is larger than the modeling noise, which will be easier to achieve in the JPEG domain rather than the spatial domain. We intend to pursue these ideas in our future work.

The code for all tested steganographic algorithms, feature extractors, the ensemble classifier, as well as BOSS-base generation scripts are available from <http://dde.binghamton.edu/download/>.

ACKNOWLEDGMENT

This work was supported by Air Force Office of Scientific 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.

APPENDIX

In this appendix, we point out that the costs of SI-UNIWARD as defined in [13] are incorrect, state how the costs are computed in the SI-UNIWARD implementation, how they should be defined, and how they are related to the costs of J-UNIWARD.

Denoting the block-wise DCT with J , the (non-rounded) DCT coefficients of the uncompressed precover image will be denoted as $J(\mathbf{P}) = \mathbf{U} \in \mathbb{R}^{n_1 \times n_2}$. The 2D array of quantized DCT coefficients is $\mathbf{X} = Q_{11}(\mathbf{U})$. Eqs. 5 and 6 in [13] define the cost of changing X_{ij} to $Y_{ij} = X_{ij} + \text{sign}(e_{ij})$ as

$$\rho_{ij}^{(\text{SI})}(\mathbf{X}) = D^{(\text{SI})}(\mathbf{X}, \mathbf{X}_{\sim ij} Y_{ij}), \quad (9)$$

where $\mathbf{X}_{\sim ij} Y_{ij}$ stands for the matrix \mathbf{X} with only X_{ij} changed to Y_{ij} and J^{-1} is the block-wise inverse DCT

(without rounding to \mathcal{I}_8). The non-additive distortion $D^{(\text{SI})}$ is defined as

$$D^{(\text{SI})}(\mathbf{X}, \mathbf{Y}) = D(\mathbf{P}, J^{-1}(\mathbf{Y})) - D(\mathbf{P}, J^{-1}(\mathbf{X})), \quad (10)$$

$$D(\mathbf{A}, \mathbf{B}) = \sum_{b=1}^3 \sum_{u,v=1}^{n_1, n_2} \frac{|W_{uv}^{(b)}(\mathbf{A}) - W_{uv}^{(b)}(\mathbf{B})|}{\sigma + |W_{uv}^{(b)}(\mathbf{A})|}, \quad (11)$$

where $W_{uv}^{(b)}(\mathbf{X})$ is the uv -th wavelet coefficient in subband b in image \mathbf{X} , σ is a stabilizing constant, and $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n_1 \times n_2}$. Using the wavelet kernel in b th subband, $\mathbf{W}^{(b)}$, the wavelet coefficients are computed using a convolution, $W_{uv}^{(b)}(\mathbf{A}) = (\mathbf{W}^{(b)} \star \mathbf{A})_{uv}$.

According to this definition, and in contrast with the claims made in [13], the costs (9) defined this way may become negative, which can be easily verified by implementing the formulas. The implementation of SI-UNIWARD available from the authors' web site uses a different formula, which always gives non-negative costs. The formula that exactly corresponds to the implementation of SI-UNIWARD should have been

$$\rho_{ij}^{(\text{SI})}(\mathbf{X}) = D^{(\text{SI})}(\mathbf{U}_{\sim ij} X_{ij}, \mathbf{U}_{\sim ij} Y_{ij}). \quad (12)$$

We now show that the costs defined this way follow the paradigm introduced in Section II, where we propose to modulate by the factor $1 - 2|e_{ij}|$ the costs of an additive scheme, which in this case is J-UNIWARD computed using the precover. Recalling that $X_{ij} = U_{ij} - e_{ij}$ and $Y_{ij} = U_{ij} + \text{sign}(e_{ij}) - e_{ij}$, we use the Dirac delta δ_{ij} to express

$$J^{-1}(\mathbf{U}_{\sim ij} X_{ij}) = J^{-1}(\mathbf{U} - \delta_{ij} e_{ij}) = \mathbf{P} + e_{ij} J^{-1}(\delta_{ij}), \quad (13)$$

$$J^{-1}(\mathbf{U}_{\sim ij} Y_{ij}) = \mathbf{P} + (\text{sign}(e_{ij}) - e_{ij}) J^{-1}(\delta_{ij}). \quad (14)$$

The linearity of convolution allows us to write (12)

$$\rho_{ij}^{(\text{SI})}(\mathbf{X}) = \sum_{b,u,v} \frac{|\text{sign}(e_{ij}) - e_{ij}| |(\mathbf{W}^{(b)} \star J^{-1}(\delta_{ij}))_{uv}|}{\sigma + |(\mathbf{W}^{(b)} \star \mathbf{P})_{uv}|} - \frac{|e_{ij}| |(\mathbf{W}^{(b)} \star J^{-1}(\delta_{ij}))_{uv}|}{\sigma + |(\mathbf{W}^{(b)} \star \mathbf{P})_{uv}|} \quad (15)$$

$$= (1 - 2|e_{ij}|) \sum_{b,u,v} \frac{|(\mathbf{W}^{(b)} \star J^{-1}(\delta_{ij}))_{uv}|}{\sigma + |(\mathbf{W}^{(b)} \star \mathbf{P})_{uv}|} \quad (16)$$

$$= (1 - 2|e_{ij}|) \rho_{ij}^{(\text{J})}(\mathbf{P}), \quad (17)$$

which is the cost of changing the ij th DCT coefficient in J-UNIWARD with the precover (rather than cover) in the denominator modulated by $1 - 2|e_{ij}|$. When the denominator uses the cover \mathbf{X} instead of the precover, we will denote the J-UNIWARD costs with $\rho_{ij}^{(\text{J})}(\mathbf{X})$.

REFERENCES

- [1] P. Bas, T. Filler, and T. Pevný. Break our steganographic system – the ins and outs of organizing BOSS. In *Proc. of the 13th IH Workshop*, volume 6958 of *LNCS*, pages 59–70, Prague, Czech Republic, May 18–20, 2011.
- [2] T. Denemark and J. Fridrich. Improving steganographic security by synchronizing the selection channel. In *3rd ACM IH&MMSec. Workshop*, Portland, Oregon, June 17–19, 2015.
- [3] T. Filler and J. Fridrich. Gibbs construction in steganography. *IEEE TIFS*, 5(4):705–720, 2010.
- [4] T. Filler, J. Judas, and J. Fridrich. Minimizing additive distortion in steganography using syndrome-trellis codes. *IEEE TIFS*, 6(3):920–935, 2011.
- [5] T. Filler, A. D. Ker, and J. Fridrich. The Square Root Law of steganographic capacity for Markov covers. In *Proc. SPIE EI*, volume 7254, pages 08 1–11, San Jose, CA, Jan. 18–21, 2009.
- [6] J. Fridrich and R. Du. Secure steganographic methods for palette images. In *Proc. of the 3rd IH Workshop*, volume 1768 of *LNCS*, pages 47–60, Dresden, Germany, Sep. 29–Oct. 1, 1999.
- [7] J. Fridrich, M. Goljan, and D. Soukal. Perturbed quantization steganography using wet paper codes. In *Proc. of the 6th ACM MMSec. Workshop*, pages 4–15, Magdeburg, Germany, September 20–21, 2004.
- [8] J. Fridrich and J. Kodovský. Rich models for steganalysis of digital images. *IEEE TIFS*, 7(3):868–882, 2011.
- [9] L. Guo, J. Ni, and Y.-Q. Shi. An efficient JPEG steganographic scheme using uniform embedding. In *Proc. of the 4th IEEE WIFS*, Tenerife, Spain, December 2–5, 2012.
- [10] L. Guo, J. Ni, and Y. Q. Shi. Uniform embedding for efficient JPEG steganography. *IEEE TIFS*, 9(5):814–825, 2014.
- [11] V. Holub and J. Fridrich. Designing steganographic distortion using directional filters. In *Proc. of the 4th IEEE WIFS*, Tenerife, Spain, December 2–5, 2012.
- [12] V. Holub and J. Fridrich. Phase-aware projection model for steganalysis of JPEG images. In *Proc. SPIE EI*, volume 9409, San Francisco, CA, February 8–12, 2015.
- [13] V. Holub, J. Fridrich, and T. Denemark. Universal distortion design for steganography in an arbitrary domain. *EURASIP Journal on Information Security*, 2014:1, 2014.
- [14] F. Huang, J. Huang, and Y.-Q. Shi. New channel selection rule for JPEG steganography. *IEEE TIFS*, 7(4):1181–1191, 2012.
- [15] A. D. Ker. A fusion of maximal likelihood and structural steganalysis. In *Proc. of the 9th IH Workshop*, volume 4567 of *LNCS*, pages 204–219, Saint Malo, France, June 11–13, 2007.
- [16] A. D. Ker, T. Pevný, J. Kodovský, and J. Fridrich. The Square Root Law of steganographic capacity. In *Proc. of the 10th ACM MMSec. Workshop*, pages 107–116, Oxford, UK, Sep. 22–23, 2008.
- [17] Y. Kim, Z. Duric, and D. Richards. Modified matrix encoding technique for minimal distortion steganography. In *Proc. of the 8th IH Workshop*, volume 4437 of *LNCS*, pages 314–327, Alexandria, VA, July 10–12, 2006.
- [18] J. Kodovský and J. Fridrich. Steganalysis of JPEG images using rich models. In *Proc. SPIE EI*, volume 8303, pages 0A 1–13, San Francisco, CA, January 23–26, 2012.
- [19] J. Kodovský, J. Fridrich, and V. Holub. Ensemble classifiers for steganalysis of digital media. *IEEE TIFS*, 7(2):432–444, 2012.
- [20] B. Li, M. Wang, and J. Huang. A new cost function for spatial image steganography. In *Proc. IEEE ICIP*, Paris, France, October 27–30, 2014.
- [21] Bin Li, Ming Wang, Xiaolong Li, Shunquan Tan, and Jiwu Huang. A strategy of clustering modification directions in spatial image steganography. *Information Forensics and Security, IEEE Transactions on*, 10(9):1905–1917, Sept 2015.
- [22] V. Sachnev, H. J. Kim, and R. Zhang. Less detectable JPEG steganography method based on heuristic optimization and BCH syndrome coding. In *Proc. of the 11th ACM MMSec. Workshop*, pages 131–140, Princeton, NJ, September 7–8, 2009.
- [23] C. Wang and J. Ni. An efficient JPEG steganographic scheme based on the block-entropy of DCT coefficients. In *Proc. of IEEE ICASSP*, Kyoto, Japan, March 25–30, 2012.
- [24] Y. Wang and P. Moulin. Perfectly secure steganography: Capacity, error exponents, and code constructions. *IEEE TIT, Special Issue on Security*, 55(6):2706–2722, 2008.