IMAGE ANONYMIZATION FOR PRNU FORENSICS : A SET THEORETIC FRAMEWORK ADDRESSING COMPRESSION RESILIENCE

Ahmed Elliethy, Gaurav Sharma

Dept. of Electrical and Computer Engineering, University of Rochester, Rochester, NY 14627
{ahmed.s.elliethy, gaurav.sharma}@rochester.edu

ABSTRACT

Image forensics using sensor photo-response nonuniformity (PRNU) provides a powerful method for associating an image with the camera that captured the image. To preserve privacy despite the availability of this powerful tool, we present a new framework for image anonymization. We formulate anonymization as a feasibility problem subject to multiple constraints that seek to ensure non-detectability of the PRNU fingerprint, visual fidelity to the original image, and compatibility to compression. A feasible anonymized image is then obtained via the method of projections onto convex sets using the inherent convexity of several constraints and convex approximations for the others. We demonstrate the effectiveness of our framework by benchmarking it over a publicly available dataset of images from multiple cameras and comparing against a recently presented alternative method. In the process we also highlight a key failing of several prior methods that fail to account for quantization in the compression process and suffer from catastrophic loss of anonymity when the anonymized image is stored in a compressed format, as would commonly be the case in realistic applications. We demonstrate specifically that the compression compatibility constraints we introduce help ensure that our method does not encounter this common pitfall.

***Index Terms*—** Image anonymization, PRNU, POCS, JPEG compression compatibility

1. INTRODUCTION

Image forensics based on the Photo Response Non-Uniformity (PRNU) of digital camera sensors has emerged as a powerful technique for several tasks such as camera source identification, integrity verification, and authentication [1–4]. While PRNU based forensics has many worthwhile applications in law-enforcement, it can also be an intrusive and undesirable invasion of privacy. For instance, whistle-blowers and journalists making images available as evidence often desire anonymity to prevent reprisals. Also individuals sharing images privately on different fora on social media may not necessarily want corporations or other individuals to be able to associate all the different images with them. For these and other reasons, tools for anonymization of images against PRNU forensics are also desirable. Such tools are the focus of the present paper.

Several counter-forensics measures have previously been proposed against PRNU forensics for both image anonymization and for false-attribution via PRNU copy attack [5, 6]. One of the common image anonymization approaches is to weaken the estimated PRNU in the original image. For example, in [7, 8], the PRNU is modeled simply as an additive noise and the image anonymization is attempted by perturbing the pixel values of the original image by subtracting an estimate of the PRNU modulated by a scaling factor that determine how much of the PRNU should be removed in order to circumvent the source camera identification detector. A recent enhancement to those methods was proposed in [9] that uses a more realistic multiplicative noise model along with a bisection search for fast and accurate estimation of the optimal PRNU scaling factor which varies across images. These approaches, while promising have the drawback that they typically disregard realistic problem considerations such as the need for the anonymized image to be stored in a compressed format. In particular, as we demonstrate in the results, for a significant class of these methods the **anonymization fails catastrophically when the images from these methods are stored in JPEG format**. Alternatively, using aggressive denoising to render the PRNU undetectable compromises visual quality. Another approach for anonymization has also been recently proposed based on seam-carving [10,11]. To maximally preserve salient content, seam-carving scales the image by sequentially removing minimum energy seams, i.e., connected horizontal or vertical paths running across the image whose removal minimally impacts the regularity of the image [12]. Because this process alters the underlying geometry of the image samples, it makes it challenging to detect the PRNU signature. However, seam carving can often distort important content in the images. In particular, the alteration of the geometry of the image is undesirable for use in computer vision and photogrammetric tasks [13] that are often applied to evidence images that journalists and whisle-blowers provide anonymously.

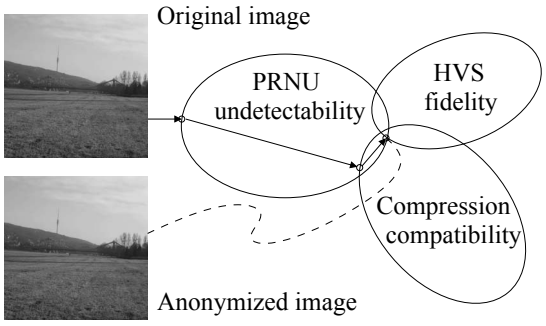

Fig. 1. Schematic of the proposed image anonymization framework. The anonymized image is obtained by imposing constraints ensuring (1) PRNU undetectability, (2) human visual system fidelity, and (3) compression compatibility. Each constraint is described exactly or approximately by one or more convex sets and the anonymized image is obtained by successive projections (POCS).

In this paper, we propose a novel framework in which image anonymization is accomplished by using set theoretic estimation [14] to synthesize a feasible image that satisfies the multiple constraints desired of the anonymized image. Specifically, we formulate constraints requiring that the anonymized image: (1) maintain undetectability with the source camera's PRNU, (2) high visual fidelity with the original image, and (3) maintain these aforementioned two attributes despite being subjected to JPEG compression, which we refer to as *compression compatibility*. Each of the constraints is represented or approximated by one or more convex

constraint sets and the anonymized image is obtained by the method of successive projections onto convex sets (POCS) [14–16]. The methodology is schematically illustrated in Fig. 1. Results obtained with our methodology on a public test image dataset demonstrate that the proposed method is indeed effective in anonymization. Importantly, we also demonstrate that **the compression compatibility constraints we propose and use in our framework overcome the catastrophic failures encountered by other methods**.

This paper is organized as follows. We introduce notational conventions and summarize PRNU based forensics in Section 2. In Section 3, we present the proposed image anonymization framework and introduce the constraints with convex approximations for the non-convex constraint sets. Then we will show our experimental results and comparison against another state of the art method in Section 4. We list our conclusions and outline future work in Section 5.

2. PRNU BASED FORENSICS

Throughout the paper, we represent 2D images (and other discrete space 2D signals) in vector form by "stacking the columns" [17]. Let $\mathbf{w} \in \mathbb{R}^{MN\times 1}$ represent the observed image with size $M \times N$, where the imaging process is modeled as [2]

$$\mathbf{w} = \mathbf{w}_0 + (\Lambda_{\mathbf{f}_0}\mathbf{w}_0) + \boldsymbol{\epsilon}, \tag{1}$$

where $\mathbf{w}_0 \in \mathbb{R}^{MN\times 1}$ is the incident light intensity, $\Lambda_{\mathbf{x}} \in \mathbb{R}^{MN\times MN}$ denotes a diagonal matrix that has the vector $\mathbf{x} \in \mathbb{R}^{MN\times 1}$ as its diagonal, $\mathbf{f}_0 \in \mathbb{R}^{MN\times 1}$ is the multiplicative PRNU pattern, and $\boldsymbol{\epsilon} \in \mathbb{R}^{MN\times 1}$ is the modeling noise.

The maximum likelihood estimate of the PRNU $\mathbf{f}_0$ from N images $\{\mathbf{w}^i\}_{i=1}^N$ that captured by the same sensor is given by [2]

$$\mathbf{f} = \Lambda_{\mathbf{s}}^{-1}\sum_{i=1}^{N}\Lambda_{\mathbf{w}^i}\mathbb{N}_{\mathbf{w}^i}\mathbf{w}^i, \tag{2}$$

where $\mathbf{s} = \sum_{i=1}^{N}\Lambda_{\mathbf{w}^i}\mathbf{w}^i$, and $\mathbb{N}_{\mathbf{w}^i}$ is the noise extraction matrix for the image $\mathbf{w}^i$. The noise extraction can be formulated in spatial, frequency, or wavelet domains. As in [2] an adaptive wavelet denoising filter [18] is used in this paper.

To decide whether a given image $\mathbf{z} \in \mathbb{R}^{NM\times 1}$ is captured from the camera used to estimate $\mathbf{f}$, first the noise $\mathbb{N}_{\mathbf{z}}\mathbf{z}$ is extracted from the image $\mathbf{z}$, then the PRNU detector uses a statistics $\rho(\mathbb{N}_{\mathbf{z}}\mathbf{z}, \mathbf{f})$ for the decision. This statistics can be correlation [6], normalized cross correlation, peak to correlation energy (PCE) [2, 19], or correlation over circular cross-correlation norm (CCN) [9, 20].

In this paper, we use the linear correlation metric, which is defined for vectors $\mathbf{x}_1$, $\mathbf{x}_2$ as $\rho(\mathbf{x}_1, \mathbf{x}_2) = \mathbf{x}_1^T\mathbf{x}_2$. We use this metric because the geometry of the image samples is unchanged in our anonymization approach, eliminating the need to address geometrical transformations such as cropping and scaling. The PRNU detector declares that the image $\mathbf{z}$ is captured with the same camera, if $\rho(\mathbb{N}_{\mathbf{z}}\mathbf{z}, \mathbf{f})$ is greater than a threshold τ.

3. PROPOSED PRNU ANONYMIZATION FRAMEWORK

Given an original image $\mathbf{x} \in \mathbb{R}^{NM\times 1}$ that is captured by a source camera whose PRNU is given by $\mathbf{f}$, our goal is to synthesize a new image $\mathbf{y} \in \mathbb{R}^{NM\times 1}$ that satisfies desired constraints. As already mentioned in the introduction, we represent the constraints, either in their original formulation or in approximate form, by convex constraint sets and obtain the anonymized image by the method of POCS. While the framework is extensible, in the present work, we specifically formulate five convex constraint sets $\mathbb{C}_1, \mathbb{C}_2, \ldots, \mathbb{C}_5$ (described subsequently) that the anonymized image must lie in, and obtain the anonymized image as the solution to the iteration

$$\mathbf{y}^{(n+1)} = P_{\mathbb{C}_5}\left(P_{\mathbb{C}_4}\left(P_{\mathbb{C}_3}\left(P_{\mathbb{C}_2}\left(P_{\mathbb{C}_1}\left(\mathbf{y}^{(n)}\right)\right)\right)\right)\right), \tag{3}$$

where $\mathbf{y}^{(n)}$ denotes the anonymized approximation at the n^{th} iteration, and

$$P_{\mathbb{C}}(\mathbf{v}) = \arg\min_{\mathbf{z}\in\mathbb{C}}\|\mathbf{v}-\mathbf{z}\|^2 \tag{4}$$

denotes the projection of the vector $\mathbf{v}$ onto the constraint set $\mathbb{C}$, i.e., the point in $\mathbb{C}$ that is closest to $\mathbf{v}$. The iteration in (3) is guaranteed to be globally convergent provided the intersection of the constraint sets is non-empty [14–16].

The projections onto the constraint sets can be analytically computed using the Kharush-Kuhn-Tucker (KKT) first order optimality conditions [21], where the Lagrange parameters need to be computed numerically. Efficient implementations of the projection operations are also obtained using either spatial, DCT, or Fourier domain implementations. Details are omitted here due to space constraints.

3.1. Original PRNU Undetectability

To decide whether a camera whose PRNU is given by $\mathbf{f}$ is used to capture a given image $\mathbf{z}$ or not, the detector checks whether the computed correlation $\rho(\mathbb{N}_{\mathbf{z}}\mathbf{z}, \mathbf{f})$ is higher than a threshold τ or not. Thus, in order to anonymize the original image, we should keep this correlation level below the PRNU detector threshold τ. We therefore represent this undetectability constraint as

$$\mathbb{C}'_1 \equiv \{\mathbf{z} : \rho(\mathbb{N}_{\mathbf{z}}\mathbf{z}, \mathbf{f}) \le \theta_d\}, \tag{5}$$

where θ_d is suitable a threshold that is less than τ, and $\mathbb{N}_{\mathbf{z}}$ is the noise extraction matrix for the image $\mathbf{z}$ as indicated earlier. We adopt the adaptive de-noising filter that is used in [2], and originally proposed in [18]. In order to incorporate this filter in our context, we describe the whole de-noising process as a series of matrix multiplications. Specifically,

$$\mathbb{N}_{\mathbf{z}} = \mathcal{W}^s\Lambda_{\sigma^{\mathbf{2}}(\mathbf{z})}\mathcal{W}^a, \tag{6}$$

where $\mathcal{W}^a$, and $\mathcal{W}^s \in \mathbb{R}^{NM\times NM}$ are the wavelet analysis and synthesis operators (represented as matrices), respectively, that are subject to the perfect reconstruction property $\mathcal{W}^s\mathcal{W}^a = \mathbf{I}$, and $\sigma^{\mathbf{2}}(\mathbf{z})$ is a scaling vector that is computed from the local variance estimate of the noise-free image, which is estimated from the image $\mathbf{z}$.

The constraint $\mathbb{C}'_1$ is nonconvex, and its convex approximation $\mathbb{C}_1$ is obtained by replacing the $\mathbb{N}_{\mathbf{z}}$ in (5) with $\mathbb{N}_{\mathbf{x}}$. Such approximation is reasonable within the feasible set because of the visual fidelity constraints (Sec. 3.2) imposed requirement for visual similarity between $\mathbf{z}$ and $\mathbf{x}$.

3.2. Anonymization imperceptibility

The anonymized image should be perceptually similar to the original one. To enforce this requirement, we propose the overall visual fidelity and the local spatial noise masking constraints, which are detailed next.

Overall visual fidelity : The human visual system (HVS) is more sensitive to lower spatial frequencies compared to higher frequencies (with an overall bandpass response that also underweights extremely low spatial frequencies). This behavior is normally modeled as a contrast sensitivity function that characterizes the relative sensitivity to each frequency [22, 23]. We therefore constrain the visual contrast sensitivity weighted difference between the original image and the anonymized one to be smaller than a threshold, which is represented by the set

$$\mathbb{C}_2 \equiv \{\mathbf{z} : \|\mathbf{H}\mathbf{z} - \mathbf{H}\mathbf{x}\|_2 \le \theta_f\}, \tag{7}$$

where θ_f is a suitable threshold, and $\mathbf{H}$ is a matrix that represents the filtering operation with a spatial linear shift invariant filter whose frequency domain representation defined as $\mathcal{H}(f_r) =$

$2.6[0.0192 + 0.114 f_r] \exp(-(0.114 f_r)^{1.1})$ represents the contrast sensitivity function from [22], where f_r is the radial frequency in cycles per degree.

Local spatial noise masking : Because the camera source identification for a given image $\mathbf{z}$ is based on the extracted noise from the image, we vary the perturbations allowed for each pixel locally based on the noise power contained in this pixel in a way that limits perturbations introduced in the noise free pixels and allows larger perturbations for the noisy pixels. This way, we selectively maintain better fidelity for noise free pixels to the original image which add a complementary visual fidelity benefit in addition to the overall one in (7), without triggering the correlation based PRNU detector used for the camera source identification task. Mathematically, we model this constraint as the set

$$\mathbb{C}_3 \equiv \{\mathbf{z} : \mathbf{l}(\mathbb{N}_\mathbf{x}\mathbf{x}) \leq |\mathbf{z} - \mathbf{x}| \leq \mathbf{u}(\mathbb{N}_\mathbf{x}\mathbf{x})\}, \tag{8}$$

where $\mathbf{l}$ and $\mathbf{u}$ represent the pixel wise lower and upper bounds respectively, and $|.|$ denotes the absolute value. We formulate $\mathbf{l}$, and $\mathbf{u}$ in a manner analogous to the texture masking constraint in [24], obtained with different motivation.

3.3. Compression compatibility

To ensure that the perturbations introduced to anonymize the original image survive storage in a compressed format, we propose successive compression decompression (SCD) PRNU undetectability and the (SCD) similarity constraints, which are detailed next.

Successive compression decompression PRNU undetectability : We model image compression in a standard transform coding framework [25], using the model of the ubiquitous JPEG compression for our description. The (SCD) operations on a given image $\mathbf{z}$ are represented as the operation

$$\mathbb{T}\{\mathbf{z}\} = \mathcal{D}^i Q[\mathcal{D}^f \mathbf{z}], \tag{9}$$

where $\mathcal{D}^f$, and $\mathcal{D}^i$ represent the forward and inverse 8×8 block discrete cosine transform (DCT) matrices respectively, and $Q[.]$ represents the quantization step, that is determined by the required quality factor and operates on coefficient by coefficient manner within the 8×8 block of the DCT transformed coefficients $\mathcal{D}^f \mathbf{z}$. We model the original PRNU undetectability after the (SCD) operations constraint as

$$\mathbb{C}'_4 \equiv \{\mathbf{z} : \rho\left(\mathbb{N}_\mathbf{z}\mathbb{T}\{\mathbf{z}\}, \mathbf{f}\right) \leq \theta_s\}, \tag{10}$$

where θ_s is a suitably chosen threshold.

Successive compression decompression similarity : To ensure that the changes introduced to anonymize an image survive compression even when the set $\mathbb{C}'_4$ is replaced with a convex approximation, we add an additional constraint that the anonymized image matches the result obtained after SCD operations, formulated as

$$\mathbb{C}'_5 \equiv \{\mathbf{z} : \mathbb{T}\{\mathbf{z}\} = \mathbf{z}\}, \tag{11}$$

where $\mathbb{T}\{.\}$ is as defined in (9).

The constraints $\mathbb{C}'_4$, and $\mathbb{C}'_5$ are nonconvex, and their convex approximation ($\mathbb{C}_4$, and $\mathbb{C}_5$) are obtained as in [24], by replacing the JPEG quantization operation with projection onto the subspace of DCT bases that are not annihilated by the quantization operation applied to the original image at a chosen quality factor.

4. EXPERIMENTAL RESULTS

To evaluate the performance of our proposed method, we conducted experiments over the Dresden Image Database (DID) [26] that contains a large set of images captured using a diverse set of camera sensors. We exclude low resolution images and, to make the computational load more tractable, crop a central 2048×2048 pixel region (so $M = N = 2048$). The resulting data set contains 14,732 images captured from 63 different camera sensors (excluding flat-field images). For each sensor, we use 50 images for the PRNU estimation via (2) and randomly select 50 images for testing PRNU based detection and anonymization, resulting in a total of $50 \times 63 = 3150$ test images. Parameters for the proposed algorithm are empirically estimated and set to $\theta_d = \theta_s = 10$, $\theta_f = MN$, for all images. For compression, we use JPEG with quality factor of 75 in all our experiments[1].

Using de-noising directly to estimate an anonymized image, typically compromises visual quality due to blurring (if the de-nosing is aggressive), or does not achieve anonymity (if the de-noising is mild). Therefore, prior approaches [7–9] for anonymization have typically proposed adaptive removal of the PRNU from the original image (in additive or multiplicative fashions) till the detection statistic falls below the detection threshold. For comparison, we use the most recent of these approaches [9] as a prototype (other methods in this class provide similar results), which we refer to as the "adaptive signature removal (**ASR**)" method. The ASR method [9] uses the CCN as a statistical measure differs slightly from the correlation measure used in our formulation (with regard to normalization and a monotonic square-root transformation). For fair and backwards compatible comparison, we report our results using CCN as in [9]. To highlight the specific impact of the compression compatibility sets we also include, in the comparisons, a de-featured version from our proposed approach called "**Proposed–**", which is obtained by dropping the compression compatibility constraints. To assess visual fidelity, we use the visual signal to noise ratio (VSNR) [28] metric.

Figure 2 shows a quantitative comparison between our proposed method and the (ASR) technique. The CCN values for the PRNU based camera sensor identification detector and the VSNR for the anonymized images obtained from the two methods are shown as scatter plots for two cases corresponding to situations where the image input to the detector is either: (1) the floating point image obtained from the algorithm with no compression (subfigure (a)) or (2) the image obtained by saving the output of the anonymization algorithm in JPEG format (subfigure (b)). We note that the second situation will be much more typical of practical use scenarios because users rarely exchange floating point images. Also, the JPEG format image can always be obtained for a second test from the uncompressed floating point representation. For orientation, it is useful to note that based on large scale tests conducted in [19], it is suggested that the detector should declare a match for CCN values[2] larger than $\sqrt{60}$, a value that gives an estimated false acceptance rate of 2.4×10^{-5}. Thus images for which the CCN is below $\sqrt{60}$ should be considered anonymized and this level is therefore indicated by a horizontal line in Figs. 2(a) and (b). Also, a large value of VSNR is desirable. From Fig. 2(a) it appears that the ASR method performs much better than the proposed method from the point of view of both visual fidelity and undetectability of the PRNU. Figure 2(b), however, illustrates that this **performance advantage of ASR over the proposed method is illusory**: once the anonymized image provided by the ASR method is stored in JPEG format, the resulting image typically gives a much high CCN value in the PRNU detector. The proposed method on the other hand yields a CCN for the PRNU detector that is much more stable under compression. Thus, **in the**

[1] The proposed methodology also works for other quality factors, although, as noted in [27], for maintaining a constant false alarm probability the threshold needs to be modified according to the quality factor.

[2] Note that the PCE = CCN^2 as shown in [20].

realistic scenario, where the anonymized image is stored in compressed JPEG format, the proposed method offers significantly better performance than the ASR method.

The ASR method introduces extremely small changes in the original image to produce the anonymized image, as is apparent from the high VSNR for this method. However, because these small changes are also substantially smaller than the quantization bins used for the JPEG compression, the anonymized images substantially revert to the original image upon being saved in JPEG format causing a catastrophic loss of the illusory anonymization. On the other hand, the compression compatibility constraint sets in the proposed method seek to ensure that the changes introduced for anonymization are preserved post-compression. Thus the changes are larger in magnitude but survive JPEG compression ensuring that a large majority of the images remain anonymized even after storage in JPEG format. The larger magnitude of the changes also implies a lower VSNR for the proposed method compared with the ASR method, which can also be seen in Figs. 2(a) and (b). The visual quality of the images for these lower VSNR values is, however, quite acceptable. A sample anonymized image obtained with the proposed method and the corresponding original image can be seen in Fig. 1 where we illustrated our algorithm (viewing of the original submitted PDF under high zoom is recommended for comparison).

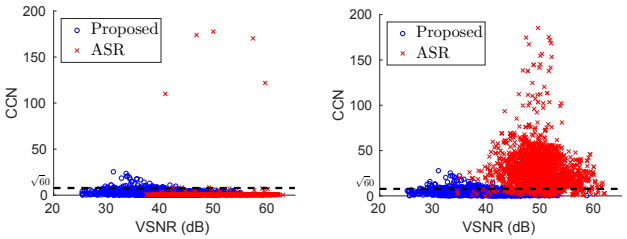

(a) Before JPEG (b) After JPEG

Fig. 2. Scatter plot of the CCN (representing PRNU detectability) versus the visual fidelity (VSNR) for anonymized images obtained using the proposed method and the ASR method [9]. Specifically, (a) represents the case when the anonymized images are retained in floating point format (Before JPEG) without any compression or file format related quantization, while (b) represent the other case when the anonymized images are JPEG compressed as would be the case for typical applications. The proposed method offers significant performance enhancement compared to the ASR method, which degrades to a rather poor level upon JPEG compression.

In Fig. 3 (a), we compare the CCN for the proposed method against the de-featured (proposed–) alternative, both for anonymized images stored in JPEG format. The figure highlights the importance of the compression compatibility sets. In the absence of the compression compatibility constraints (i.e. for proposed–), the anonymized images have a high CCN compared to the proposed method. In Fig. 3(b) we present detection error trade-off (DET) curves [29] for comparing the performance of camera sensor identification detectors that operate by thresholding the CCN for the proposed method, for proposed–, and for ASR (all for images stored in JPEG format). The DET curve for original non-anonymized images is also included in the figure for reference. The detector has almost an identical performance on the original images and the anonymized images obtained using the ASR method after JPEG compression, while our method's anonymized images significantly degrade the performance of the detector. The de-featured (proposed–) method performs better than the ASR method, but is markedly worse than the proposed method.

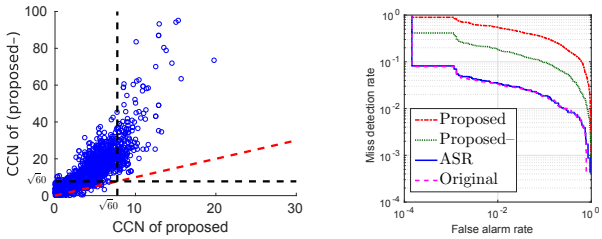

(a) (b)

Fig. 3. A comparison of the CCN (PRNU detectability) of our proposed method and the (proposed–) method is shown in (a), while we show in (b) the performance of the PRNU detector in terms of detection error trade-off (DET) curve [29] when presented the original images, and when presented the anonymized images obtained using our proposed method, the (proposed–) method, and the ASR method after these anonymized images are JPEG compressed.

Table 1 shows the percentage of the anonymized images that circumvent the PRNU based camera identification when the CCN threshold is set at $\sqrt{60}$ (as justified earlier). The numerical values in the table reinforce the conclusions drawn from Figs. 2 and 3. Our proposed method manages to significantly degrade the detector performance compared with the ASR, and the (proposed–) methods, thereby offering much better privacy.

	ASR	Proposed–	Proposed
After JPEG	9.09%	42.29%	**91.1**%
Before JPEG	**99.65**%	94.04%	97.45%

Table 1. The percentage of the anonymized images for which the PRNU detector fails to make the correct association. Results are reported for each of the methods compared for the situations when the anonymized image is JPEG compressed and when the image is retained in floating point format (Before JPEG) without any compression or file format related quantization. Note the proposed method offers very similar performance for the two cases whereas the ASR method's performance degrades to a rather poor level upon JPEG compression.

5. CONCLUSION

In this paper, we propose a new framework for image anonymization against PRNU forensics. Instead of an explicit method for computing the anonymized image, our framework obtains the image implicitly via an image feasibility problem with constraints that impose undetectability, visual fidelity, and compression compatibility. Convex representations/approximations of these constraints allow computation of the anonymized image using the method of POCS. Results obtained with our method and compared against a recently proposed alternative that is representative of the majority of prior approaches highlight that the framework is not only effective at anonymization but also the compression compatibility constraints that we introduce fix the catastrophic compromise of anonymity against PRNU forensics when images are stored in JPEG format, as would commonly be the case in actual practice. Highlighting this failure is also a contribution of the present work.

The framework we propose is flexible and extensible. Through the addition, deletion, and substitution of constraints we can address additional or different requirements. To demonstrate this flexibility, in future work, we plan to develop extensions that address additional counter-forensics requirements beyond anonymization, for instance, misdirection of attribution.

6. REFERENCES

[1] M. Chen, J. Fridrich, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE Trans. Info. Forensics and Security*, vol. 3, no. 1, pp. 74–90, March 2008.

[2] J. Fridrich, "Digital image forensics," *IEEE Sig. Proc. Mag.*, vol. 26, no. 2, pp. 26–37, March 2009.

[3] H. Cao and A. Kot, "Manipulation detection on image patches using FusionBoost," *IEEE Trans. Info. Forensics and Security*, vol. 7, no. 3, pp. 992–1002, June 2012.

[4] V. Conotter, P. Comesana, and F. Perez-Gonzalez, "Forensic detection of processing operator chains: Recovering the history of filtered JPEG images," *IEEE Trans. Info. Forensics and Security*, vol. 10, no. 11, pp. 2257–2269, Nov 2015.

[5] R. Böhme and M. Kirchner, "Counter-forensics: Attacking image forensics," in *Digital Image Forensics*. Springer, 2013, pp. 327–366.

[6] M. Steinebach, H. Liu, P. Fan, and S. Katzenbeisser, "Cell phone camera ballistics: attacks and countermeasures," in *Proc. SPIE: Multimedia on Mobile Devices*, vol. 7528, Jan. 2010.

[7] C.-T. Li, C.-Y. Chang, and Y. Li, "On the repudiability of device identification and image integrity verification using sensor pattern noise," in *Information Security and Digital Forensics*. Springer, 2010, pp. 19–25.

[8] A. Karaküçük and A. E. Dirik, "Adaptive photo-response nonuniformity noise removal against image source attribution," *Digital Investigation*, vol. 12, pp. 66–76, 2015.

[9] H. Zeng, J. Chen, X. Kang, and W. Zeng, "Removing camera fingerprint to disguise photograph source," in *IEEE Intl. Conf. Image Proc.*, Sept 2015, pp. 1687–1691.

[10] S. Bayram, H. Sencar, and N. Memon, "Seam-carving based anonymization against image & video source attribution," in *IEEE Intl. Workshop on Multimedia Sig. Proc.*, Sept 2013, pp. 272–277.

[11] A. Dirik, H. Sencar, and N. Memon, "Analysis of seam-carving-based anonymization of images against PRNU noise pattern-based source attribution," *IEEE Trans. Info. Forensics and Security*, vol. 9, no. 12, pp. 2277–2290, Dec 2014.

[12] S. Avidan and A. Shamir, "Seam carving for content-aware image resizing," *ACM Trans. Graph.*, vol. 26, no. 3, Jul. 2007.

[13] C. Fillion and G. Sharma, "Detecting content adaptive scaling of images for forensic applications," in *Proc. SPIE: Media Forensics and Security XII*, vol. 7541, Jan. 2010, pp. 7541Z1–12.

[14] P. Combettes, "The foundations of set theoretic estimation," *Proc. IEEE*, vol. 81, no. 2, pp. 182–208, Feb 1993.

[15] L. M. Bregman, "The method of successive projection for finding a common point of convex sets," *Dokl. Akad. Nauk. USSR*, vol. 162, no. 3, pp. 487–490, 1965.

[16] L. G. Gubin, B. T. Polyak, and E. T. Raik, "The method of projections for finding the common point of convex sets," *USSR Comput. Math. and Phys.*, vol. 7, no. 6, pp. 1–24, 1967.

[17] R. C. Gonzales and P. Wintz, *Digital Image Processing*, 2nd ed. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1987.

[18] M. Mıhçak, I. Kozintsev, and K. Ramchandran, "Spatially adaptive statistical modeling of wavelet image coefficients and its application to denoising," in *IEEE Intl. Conf. Acoust., Speech, and Signal Proc.*, vol. 6, Mar 1999, pp. 3253–3256.

[19] M. Goljan, J. Fridrich, and T. Filler, "Large scale test of sensor fingerprint camera identification," in *Proc. SPIE: Media Forensics and Security XI*, vol. 7254, Jan. 2009.

[20] X. Kang, Y. Li, Z. Qu, and J. Huang, "Enhancing source camera identification performance with a camera reference phase sensor pattern noise," *IEEE Trans. Info. Forensics and Security*, vol. 7, no. 2, pp. 393–402, April 2012.

[21] D. G. Luenberger and Y. Ye, *Linear and Nonlinear Programming*, 4th ed. New York, NY: Springer Publishing Co., Inc., 2015.

[22] J. Mannos and D. Sakrison, "The effects of a visual fidelity criterion on the encoding of images," *IEEE Trans. Info. Theory*, vol. 20, no. 4, pp. 525–536, Jul 1974.

[23] B. A. Wandell, *Foundations of Vision*. Sunderland, MA: Sinauer Associates, Inc., 1995.

[24] H. Altun, A. Orsdemir, G. Sharma, and M. Bocko, "Optimal spread spectrum watermark embedding via a multistep feasibility formulation," *IEEE Trans. Image Proc.*, vol. 18, no. 2, pp. 371–387, Feb 2009.

[25] "Special issue on transform coding: Past, present, and future," *IEEE Sig. Proc. Mag.*, vol. 18, no. 5, Sep. 2001.

[26] T. Gloe and R. Böhme, "The Dresden image database for benchmarking digital image forensics," *Journal of Digital Forensic Practice*, vol. 3, no. 2-4, pp. 150–159, 2010.

[27] M. Goljan, M. Chen, P. Comesaña, and J. Fridrich, "Effect of compression on sensor-fingerprint based camera identification," in *IS&T Electronic Imaging: Media Watermarking, Security, and Forensics*, 2016, San Francisco, CA, 14-18 February 2016.

[28] D. Chandler and S. Hemami, "VSNR: A wavelet-based visual signal-to-noise ratio for natural images," *IEEE Trans. Image Proc.*, vol. 16, no. 9, pp. 2284–2298, Sept 2007.

[29] A. Martin, G. Doddington, T. Kamm, M. Ordowski, and M. Przybocki, "The DET curve in assessment of detection task performance," in *Proc. Eurospeech Conf.*, 1997, pp. 1895–1898.