

Beyond Non-Unique Artifacts: SVD-Based PRNU Recovery for Samsung Device Identification

David Vázquez-Padín
dvazquez@gts.uvigo.es
atlanTTic, Universidade de Vigo
Vigo, Spain

Fernando Pérez-González
fperez@gts.uvigo.es
atlanTTic, Universidade de Vigo
Vigo, Spain

Abstract

We investigate non-unique artifacts in images captured by Samsung smartphones and their impact on PRNU-based source camera verification. These artifacts, linked to Exynos-equipped devices, manifest as 6 distinct periodic patterns identified across models released between 2018 and 2025, and lead to fingerprint collisions that degrade verification performance. For the most prevalent pattern, we restore reliable verification by applying SVD to separate block-induced artifacts from the underlying PRNU. We further exploit the regularity of these artifacts to estimate and compensate for HDR-induced local misalignments, enabling fingerprint synchronization in HDR images. Experimental results show that combining SVD-based PRNU recovery with HDR-aware synchronization significantly improves detection performance, in several cases approaching perfect verification. While newer artifact patterns remain more challenging, this work constitutes a first step toward mitigating fingerprint collisions in Samsung devices.

CCS Concepts

• Computing methodologies → Image processing.

Keywords

Image forensics, SVD, Source Camera Verification, PRNU, Fingerprint Collision, Samsung

ACM Reference Format:

David Vázquez-Padín and Fernando Pérez-González. 2026. Beyond Non-Unique Artifacts: SVD-Based PRNU Recovery for Samsung Device Identification. In *ACM Workshop on Information Hiding and Multimedia Security (IH&MMSec '26)*, June 17–19, 2026, Firenze, Italy. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3785353.3815090>

1 Introduction

In image forensics, the de facto standard for verifying whether an image originates from a given camera is the sensor pattern noise known as Photo-Response Non-Uniformity (PRNU). Although PRNU has been successfully applied in forensic investigations, its reliability is increasingly challenged by advances in computational photography. Early inconsistencies stemmed from in-camera operations such as distortion correction [20] and High Dynamic Range (HDR) processing [11], which required specialized compensation

techniques. More recently, a critical issue has emerged in modern smartphones: the loss of PRNU uniqueness [14]. This results from sophisticated computational photography and AI-based enhancements, which introduce Non-Unique Artifacts (NUAs) that can cause false positive attributions. These phenomena undermine the assumption of the PRNU as a unique identifier and raise significant concerns about its continued role as the forensic standard [16].

Following the work of Iuliani *et al.* [14], which questioned the long-assumed uniqueness of PRNU, several studies have investigated the origins of the reported fingerprint collisions. For Digital Single-Lens Reflex (DSLR) cameras, Butora and Bas [9] showed that many false positives stem from software-specific artifacts introduced by Adobe during raw image development. In the smartphone domain, Baracchi *et al.* [7] initiated the study of NUAs, focusing on the synthetic pattern embedded in Apple portrait-mode images, later further characterized in [22]. Around the same time, Albisani *et al.* [6] reported NUAs in Samsung, Huawei, and Apple devices, while Bhat and Bianchi [21] investigated similar artifacts in Samsung, Huawei, and Xiaomi smartphones. These works enable practitioners to identify NUA-affected images, either excluding them from reference fingerprint estimation or lowering confidence in attribution results. More recently, Montibeller *et al.* [19] showed that visible watermarks can lead to false positives in Xiaomi devices and observed that spikes in the Discrete Fourier Transform (DFT) magnitude are correlated with the absence of NUAs responsible for false positives in Samsung A50 devices.

While identifying NUAs is a necessary step to prevent fingerprint collisions, Albisani *et al.* [6] reported that 100% of the analyzed flatfield blocks from Samsung S9 and S9+ devices were mutually correlated, preventing the extraction of a valid PRNU fingerprint. This observation highlights that merely detecting NUAs is insufficient; their impact must be mitigated to restore reliable source attribution. In this work, we first clarify the distinct sources of artifacts responsible for the high correlations observed in [6] and contrast them with those reported in [21]. We then propose a subspace-decomposition strategy to isolate and suppress the most prevalent NUA-induced pattern identified in Samsung devices, thereby enabling recovery of the underlying sensor PRNU. To the best of our knowledge, this is the first study to demonstrate effective suppression of device-induced periodic NUAs using an SVD-based subspace separation framework for source camera verification. Our main contributions are summarized as follows:

- Identification of 6 distinct periodic patterns in images captured by Exynos-equipped Samsung smartphones (released between 2018 and 2025), which trigger false positives in PRNU-based source camera verification.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

IH&MMSec '26, Firenze, Italy

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2376-6/2026/06

<https://doi.org/10.1145/3785353.3815090>

- Development of an SVD-based subspace separation method to recover the underlying sensor PRNU by suppressing periodic artifacts.
- Introduction of an HDR-aware fingerprint synchronization strategy to estimate and compensate for local HDR-induced misalignments.
- Experimental validation demonstrating that the combination of SVD-based PRNU recovery and HDR synchronization substantially improves verification performance for the most prevalent periodic pattern.

The remainder of the paper is organized as follows. Sect. 2 reviews the fundamentals of PRNU-based source camera verification. Sect. 3 characterizes the periodic artifacts observed in Exynos-equipped devices. Sect. 4 presents the proposed SVD-based PRNU recovery method (Sect. 4.1) and the HDR-aware synchronization strategy (Sect. 4.2). Experimental results are reported in Sect. 5, and conclusions and future research directions are discussed in Sect. 6.

2 PRNU-based Source Camera Verification

Lukáš *et al.* [17] introduced the use of the PRNU, a sensor-specific noise pattern, for verifying whether a given image was captured by a specific camera. The assumed model of the sensor output for an $H \times W$ single-channel image $\mathbf{Y}$ is $\mathbf{Y} = (\mathbf{1} + \mathbf{K}) \circ \mathbf{X} + \mathbf{\Theta}$, where $\mathbf{K}$ is the PRNU, $\mathbf{X}$ the scene intensity, and $\mathbf{\Theta}$ other noise sources [10]. Based on this model, an estimate of a camera's PRNU fingerprint $\hat{\mathbf{K}}$ can be obtained from L images $\{\mathbf{Y}_l\}_{l=1}^L$ by first denoising each image, i.e., $F(\mathbf{Y}_l)$ (where the filtering $F(\cdot)$ comes from [15]), and computing residues $\mathbf{W}_l = \mathbf{Y}_l - F(\mathbf{Y}_l)$. In the general case, where NUAs may be present in the reference images, we denote the resulting estimate as $\hat{\mathbf{F}}$, computed following [10] as

$$\hat{\mathbf{F}} = \left(\sum_{l=1}^L \mathbf{W}_l \circ \mathbf{Y}_l \right) \circ \left(\sum_{l=1}^L \mathbf{Y}_l \circ \mathbf{Y}_l \right)^{\circ-1}, \quad (1)$$

where $\circ$ denotes the Hadamard product and $(\cdot)^{\circ-1}$ the element-wise inverse. The notation $\hat{\mathbf{F}}$ emphasizes that this estimate may include both the sensor PRNU and additional NUA-induced patterns. In the absence of NUAs, $\hat{\mathbf{F}}$ reduces to the conventional PRNU estimate $\hat{\mathbf{K}}$.

Now, given a test image $\mathbf{Y}$ with residue $\mathbf{W} = \mathbf{Y} - F(\mathbf{Y})$ and a PRNU estimate $\hat{\mathbf{K}}$, the verification of the camera source can be framed as a binary hypothesis test. The null hypothesis, $\mathcal{H}_0$, assumes that $\mathbf{W}$ does not contain the PRNU $\mathbf{K}$, while the alternative, $\mathcal{H}_1$, assumes that it does. If $\mathcal{H}_1$ holds, the test image $\mathbf{Y}$ likely originates from the camera with PRNU $\mathbf{K}$. In practice, the hypothesis test is evaluated using the Peak-to-Correlation Energy (PCE) [13], which is defined as follows

$$\text{PCE}(\mathbf{W}, \hat{\mathbf{K}} \circ \mathbf{Y}) \triangleq \frac{(H \cdot W - |\mathcal{N}|) \cdot (\rho_{(0,0)}(\mathbf{W}, \hat{\mathbf{K}} \circ \mathbf{Y}))^2}{\sum_{(s_1, s_2) \in \mathcal{S} \setminus \mathcal{N}} (\rho_{(s_1, s_2)}(\mathbf{W}, \hat{\mathbf{K}} \circ \mathbf{Y}))^2}, \quad (2)$$

where $\rho_{(s_1, s_2)}(\mathbf{A}, \mathbf{B})$ denotes the Normalized Cross-Correlation (NCC) between $\mathbf{A}$ (with sample mean $\bar{\mathbf{A}}$) and $\mathbf{B}$ (with sample mean $\bar{\mathbf{B}}$) for a given shift (s_1, s_2) , i.e.,

$$\rho_{(s_1, s_2)}(\mathbf{A}, \mathbf{B}) \triangleq \frac{\sum_{i,j} (A_{i,j} - \bar{\mathbf{A}}) \cdot (B_{i+s_1, j+s_2} - \bar{\mathbf{B}})}{(\sum_{i,j} (A_{i,j} - \bar{\mathbf{A}})^2 \sum_{i,j} (B_{i+s_1, j+s_2} - \bar{\mathbf{B}})^2)^{\frac{1}{2}}}.$$

In (2), $\mathcal{S}$ denotes the set of all possible spatial shifts, and $\mathcal{N}$ is a small 11×11 neighborhood around the origin $(0, 0)$, where the NCC

peak is expected under the assumption of no cropping, which holds in our experimental setup. For readability, the zero-shift correlation $\rho_{(0,0)}(\mathbf{A}, \mathbf{B})$ is denoted hereafter as $\rho(\mathbf{A}, \mathbf{B})$.

The final decision is obtained by comparing the PCE value to a predefined threshold τ :

$$\text{PCE}(\mathbf{W}, \hat{\mathbf{K}} \circ \mathbf{Y}) \underset{\mathcal{H}_0}{\overset{\mathcal{H}_1}{\gtrless}} \tau. \quad (3)$$

In practice, a threshold $\tau = 60$ is typically adopted, as validated in [13] on over a million images from 6,896 cameras, achieving a false alarm rate of 10^{-5} and a detection rate of 97.6%.

3 Samsung Image Artifacts

Previous studies by Albisani *et al.* [6] and Bhat and Bianchi [21] showed that NUAs present in images captured by Samsung devices (primarily from the Galaxy S series) induce false positives in PRNU-based source camera verification. Later, Montibeller *et al.* [19] reported that spikes in the Fourier domain appear to correlate with the absence of the NUAs responsible for false positives in Samsung A50 devices. Since the artifacts identified in these works differ in nature, we were motivated to investigate the origin of Samsung-related NUAs that exhibit stronger correlations than the underlying PRNU.

A key difference between the datasets used in [6] and [21] lies in the specific device variants analyzed. Although similar Samsung models were considered (Galaxy S9/S9+ in the former and S10/S10+ in the latter), the regional versions differed. In particular, the model codenames ended in “F” in [6] and in “U” in [21]. Further inspection revealed that the former dataset consisted of images captured with the *Global versions* of the S9 and S9+ equipped with Exynos chipsets, whereas the latter relied on the *USA versions* of the S10 and S10+ powered by Qualcomm Snapdragon chipsets. Since different chipsets typically employ a different Image Signal Processor (ISP), this distinction likely explains the discrepancies in the observed forensic traces.

To further analyze these differences, we collected a subset of images from the same Samsung models (Galaxy S9/S9+ and S10/S10+) but different regional variants, namely the USA versions equipped with Snapdragon chipsets and their corresponding Global counterparts incorporating Exynos chipsets. As summarized in Tab. 2 (see Sect. 5), images from the USA variants were drawn from the dataset in [14] (also used in [21]), where multiple units of the same model are grouped by device type and version without individual identifiers. On the other hand, images from the Global variants (devices SD01-SD08) were obtained from public datasets such as [6] and FloreView [8], publicly available out-of-camera samples from GSMArena reviews [4], as well as from our own devices.¹

After assembling this subset of images, we computed a reference estimate $\hat{\mathbf{F}}$ from one of our Exynos-equipped devices, specifically a Samsung Galaxy S10 (device ID SD06 in Tab. 2). The estimate was obtained according to (1) using $L = 20$ flat-field images captured in default photo mode. We then evaluated the PCE, as defined in (2), between this reference estimate $\hat{\mathbf{F}}$ (which very likely contains NUAs) and the residuals $\mathbf{W}$ of all S9, S9+, S10, and S10+ test images. The

¹The dataset is available at <https://github.com/dvazquezpadin/samsung-NUAs>





Figure 1: Scatter plots of PCE values obtained with default-mode reference estimates. (a) PCE between the S10 (SD06) estimate $\hat{F}$ and S9/S9+/S10/S10+ test images, grouped by device version. (b) PCE for Galaxy A series devices (AD06-AD13), with positives and negatives shown in green and red. Values largely exceeding $\tau = 60$ indicate the presence of NUAs in $\hat{F}$.

results, grouped by device version, are shown in the scatter plot of Fig. 1(a).

The scatter plot reveals two main findings. First, the NUAs present in the Exynos/Global variants are uncorrelated with those observed in the Snapdragon/USA versions. This distinction is important because the studies in [14] and [21] analyzed only USA models; consequently, their conclusions apply exclusively to that regional variant.² Second, within the Exynos/Global variants, it is not possible to distinguish between individual devices or models, as most PCE values exceed the commonly adopted threshold of $\tau = 60$. This behavior is consistent with previous observations for the S9 (SD01) and S9+ (SD03) in [6], and later for the S10 (SD05) and S10+ (SD08) in [8]. Furthermore, the few PCE values from these devices that fall below the threshold correspond to images processed through a Multi-Frame Processing (MFP) pipeline (e.g., HDR rendering from stacked frames), and are labeled accordingly in Fig. 1(a).³ Misdetectors arising from HDR-induced PRNU misalignments are discussed in Sect. 4.2.

The collisions observed among the considered Exynos-equipped devices can be attributed to a shared characteristic in images captured under the default photo mode: the autocorrelation of their residues exhibits a diagonal periodic pattern (whose magnitude exceeds that of the PRNU signal) that aligns with the autocorrelation of the fingerprint extracted from device SD06 (see Fig. 3(a)). To emphasize this structure, we restrict the visualization to the central 801×801 region and mark the strongest positive peaks with Δ and the strongest negative peaks with ∇ . Within this region, prominent positive peaks appear, for instance, at integer multiples of the coordinates (60, 65), suggesting the presence of block-based




Figure 2: Examples from the PrnuModernDevices dataset [6] (CC BY-SA 4.0) showing Samsung diagonal artifacts in images captured with Samsung Galaxy S9 (a) and S9+ (c). The corresponding block-wise NCC maps between the S10 (SD06) estimate $\hat{F}$ and image residues W are shown in (b) and (d).

processing. Although the exact origin of the observed periodicity remains unclear, evidence suggests that it does not stem from intrinsic sensor non-uniformities. Rather, it is likely introduced within the ISP pipeline, possibly due to an ordered dithering operation applied during bit-depth reduction (e.g., from 12-16-bit RAW to 8-bit JPEG), as in devices such as the Galaxy S10, to mitigate color banding artifacts.

Consistent with this hypothesis, the correlation with the NUA-affected reference estimate $\hat{F}$ is more pronounced in spatially uniform regions, where banding-suppression mechanisms are expected to be applied. For illustration, we computed the block-wise (21×21) NCC map (smoothed with a 5×5 box filter) between the S10 (SD06) reference estimate $\hat{F}$ and residues from S9 (SD01) and S9+ (SD03) images used in [6]. While NUAs are present across almost the entire frame, higher correlation values are observed in flat regions (e.g., sky in Fig. 2(b); and walls or floors in Fig. 2(d)), whereas textured areas (e.g., trees in Fig. 2(b)) exhibit weaker responses. More recent

²In this work, we do not further investigate NUAs in Snapdragon/USA models, leaving such analysis to future research.

³MFP labeling was performed by searching the strings MHDR, LHDR, and MFP3 in the JPEG APP4 marker containing proprietary Samsung metadata.

devices incorporate advanced sensor technologies (e.g., ISOCELL with TETRAPIXEL or NONAPIXEL architectures [5]) together with increasingly sophisticated Neural Processing Units (NPUs). These evolutions introduce new artifact structures, as detailed below, whose nature differs from those observed in the S9 and S10 series and is likely to become progressively more complex to characterize, as further discussed in Sect. 5.2.

Returning to the diagonal correlation pattern shown in Fig. 3(a), positive peaks along the diagonal were previously reported by Albisani *et al.* [6, Fig. 4] for the S9 and S9+ devices. Our analysis extends these findings by revealing the presence of periodic negative correlations and demonstrating that these artifacts are not limited to Exynos-equipped S9 and S9+ models, but also appear in S10 and S10+ devices. After establishing that the diagonal artifacts are very likely associated with the ISP of Exynos chipsets, we expanded the dataset to include images from as many Exynos-equipped devices as possible in order to determine whether the same or different periodic patterns occur across Samsung models. This extended dataset incorporates additional samples from academic sources such as ForensiCam-215k [12], as well as publicly available out-of-camera images from Flickr [3] and GSMArena reviews [4], together with images acquired using our own devices, as summarized in Tab. 2. Although we aimed to include all Exynos-based devices, some models could not be analyzed due to the lack of publicly available samples (e.g., S21 FE Version 2). Moreover, devices such as the S20 Ultra 5G and S21 Ultra 5G were excluded because no NUAs were detected at their native 12MP outputs, despite evidence at full resolution (108MP); this case is left for future work.

The broader analysis revealed that additional periodic structures are present beyond the pattern initially observed in the Galaxy S9/S9+ and S10/S10+. To systematically identify and categorize these artifacts, we analyzed all Exynos-equipped devices summarized in Tab. 2 under default photo mode and grouped them according to their autocorrelation patterns. The resulting classification is summarized in Tab. 1, which lists the chipset generation, sensor and image resolutions, release year, and the associated number of the periodic artifact pattern for each device.

Overall, six distinct diagonal correlation patterns (Patterns 1-6) were identified across Exynos-equipped devices released between 2018 and 2025. These patterns correspond to different autocorrelation structures, illustrated in Fig. 3, and are shared across multiple models and chipset generations. Pattern 1 is the most prevalent and primarily observed in early Galaxy S and Note models (S9/S9+/S10/S10+/Note variants), whereas subsequent Exynos platforms introduce progressively different periodic structures (Patterns 2-6), particularly in later A series and more recent S series devices. Interestingly, Pattern 2 is also observed in non-Samsung smartphones (e.g., vivo models) that integrate Exynos chipsets, suggesting that these NUAs are tied to the chipset/ISP pipeline rather than to the device brand itself. In some cases, minor variations arise due to resolution changes or in-camera resizing operations, as indicated in Tab. 1. This systematic grouping shows that fingerprint collisions are not device-specific but pattern-dependent, with each pattern exhibiting a characteristic autocorrelation signature.

As an illustrative example, we examined Galaxy A series devices AD06-AD13 (Tab. 2) by analyzing default photo mode images from




Figure 3: Autocorrelation of the reference estimates $\hat{F}$ from representative Exynos-equipped devices: S10 (SD06) (a), vivo X60 Pro (VD02) (b), A33 5G (AD03) (c), A53 5G (AD06) (d), A25 5G (AD01) (e), and A54 5G (AD08) (f), illustrating the six identified NUA patterns (Patterns 1-6) responsible for fingerprint collisions.

two units per model. Reference estimates $\hat{F}$ were extracted following the procedure in Sect. 2, using reference images from devices AD06 (A53 5G), AD08 (A54 5G), AD10 (A55 5G), and AD12 (A56 5G). The resulting PCE values, grouped by device version, are shown in Fig. 1(b), where positive and negative matches are depicted in green and red, respectively. Most values (regardless of match class) cluster above 10^4 , indicating the presence of strong NUAs similar to those observed in the Galaxy S series. As in the previous analysis, different marker border colors identify samples obtained from MFP-processed images (e.g., HDR) and digitally zoomed images (as determined from EXIF metadata), since these conditions may yield PCE values close to or below the decision threshold if not corrected. In terms of pattern attribution, the reference estimate of the A53 5G (AD06) yields Pattern 4 (see Fig. 3(d)), whereas the estimates of the A54 5G (AD08) and A55 5G (AD10) exhibit NUAs consistent with Pattern 6 (Fig. 3(f)). The A56 5G (AD12) case is omitted for brevity, as its NUA-induced pattern matches Pattern 6 once resized from 4000×3000 to 4080×3060 .

4 SVD-Based Mitigation of Fingerprint Collisions

So far, we have shown that capturing reference images in the default photo mode on the Exynos-equipped devices listed in Tab. 1 leads to fingerprint collisions, making PRNU-based source camera verification ineffective. This naturally raises the question of whether the PRNU remains present in the captured images, albeit at lower power. Experiments conducted with our Galaxy S10 (SD06) indicate that storing images in HEIF format (instead of JPEG) does not mitigate the issue, as the reported Pattern 1 persists and collisions remain. In contrast, capturing reference images in PRO mode produces raw images (in DNG format) whose residues do not exhibit NUAs.⁴

⁴Raw images in DNG format were converted to uncompressed RGB images using dcraw [2] with default settings.

Table 1: Classification of Exynos-equipped devices according to their NUA pattern. The table reports chipset model number, sensor and image resolutions, smartphone model with its release year, and the associated Pattern label (see Fig. 3).

Exynos Model	Sensor Res.	Image Res. ($H \times W$)	Smartphone Model	Release Year	Pattern Label
9810	12MP	4032x3024	S9	2018	1
			S9+	2018	
			Note9	2018	
			Note10 Lite	2020	
9820	12MP	4032x3024	S10	2019	1
			S10+	2019	
9825	12MP	4032x3024	Note10	2019	1
			Note10+	2019	
980	48MP	4000x3000	A51 5G	2020	2
990	12MP	4032x3024	S20	2020	1
	108MP	4000x3000	S20 Ultra	2020	1
1080	48MP	4000x3000	vivo X60 [‡]	2021	2
			vivo X60 Pro [‡]	2021	
2100	12MP	4032x3024	S21	2021	1*
			S21+	2021	
2200	108MP	4000x3000	S22 Ultra	2022	2
	48MP	4000x3000	A33 5G	2022	3
1280	64MP	4624x3468	A53 5G	2022	4
	50MP	4080x3060	A25 5G	2023	5
1380	50MP	4080x3060	A54 5G	2023	6
			A35 5G	2024	
			A26 5G	2025	
1480	50MP	4080x3060	A55 5G	2024	6
2400e	50MP	4000x3000	S24 FE	2024	6†
			S24	2024	
			S24+	2024	
			Z Flip7 FE	2025	
2400	50MP	4000x3000	S25 FE	2025	6†
			S25 FE	2025	
1580	50MP	4000x3000	A56 5G	2025	6†
2500	50MP	4000x3000	Z Flip7	2025	6†

[‡] Only Chinese market versions.

* The pattern is only present under the HDR mode.

† Pattern 6† matches Pattern 6 after resolution alignment via resizing.

Building on this observation, we evaluated whether an estimate $\hat{F}_{\text{RAW}}$ obtained according to (1) from $L = 20$ flatfield PRO-mode raw images of the S10 (SD06) wide sensor can correctly identify default photo mode images captured with the same wide sensor (positive class) while rejecting images acquired with the SD06 telephoto sensor (negative class), which has the same resolution but a different PRNU. The evaluation set consisted of 71 default photo mode images captured with the same wide sensor (positives) and 40 default photo mode images captured with the SD06 telephoto sensor (negatives).

The resulting ROC curve is shown in Fig. 4, together with the curve obtained using the default-mode reference estimate $\hat{F}$, computed according to (1) from $L = 20$ flatfield JPEG images captured in default photo mode with the SD06 wide sensor. Because Pattern 1 artifacts are embedded in images captured in default photo mode by both the wide and telephoto sensors, the resulting estimate $\hat{F}$ (dashed line) is affected by these NUAs, leading to poor discrimination performance. In contrast, the PRO-mode raw reference




Figure 4: ROC curves for discrimination between the wide and telephoto sensors of the S10 device (SD06). The default-mode reference estimate $\hat{F}$ (dashed), affected by Pattern 1 NUAs, is compared with the PRO-mode raw estimate $\hat{F}_{\text{RAW}}$ (solid), from which $\hat{K}$ is constructed. The raw-based estimate restores sensor discrimination, while the remaining errors are due to HDR-induced misalignments.

estimate $\hat{F}_{\text{RAW}}$ (solid line) substantially improves sensor discrimination, confirming that the underlying PRNU remains present in default photo mode images and that the true sensor fingerprint K is well approximated by $\hat{F}_{\text{RAW}}$. Since raw images are not affected by NUAs, we construct the PRNU estimate as $\hat{K} = \hat{F}_{\text{RAW}}$. Perfect detection is nevertheless not achieved due to the presence of HDR images in the evaluation set. As discussed in Sect. 4.2, compensating for HDR-induced misalignments is necessary to attain perfect verification performance.

These results might suggest, at first glance, that employing the PRO-mode raw fingerprint $\hat{K}$ could serve as a practical solution to mitigate NUA-induced collisions. In practice, however, this is not feasible in forensic scenarios, as it requires access to the original device to acquire raw images. Moreover, PRO mode does not provide DNG output for all sensors (e.g., the S10 telephoto sensor), and many mid-range devices (such as Galaxy A models) do not support raw capture at all, offering only JPEG images with the embedded NUAs.

These limitations motivate the need for alternative strategies capable of suppressing the periodic artifacts directly from the reference estimate $\hat{F}$ while preserving the underlying PRNU. In the following, we address this challenge by employing an SVD-based subspace decomposition to separate NUA-induced dependencies from the sensor PRNU within a signal-noise framework.

4.1 SVD-Based PRNU Recovery via Subspace Separation

Since the PRNU persists in images acquired under default photo mode, the estimate $\hat{F}$, computed from L NUA-affected images according to (1), can be modeled as:

$$\hat{F} = K + P + \Phi,$$




Figure 5: NCC between the top-left 65×60 block of the S10 (SD06) reference estimate $\hat{\mathbf{F}}$ (containing Pattern 1 NUAs) and the remaining samples of $\hat{\mathbf{F}}$. The highest-magnitude peaks are highlighted with white markers for clarity. The main diagonal is emphasized with circled markers due to its relevance in the subsequent SVD-based subspace analysis.

where $\mathbf{K}$ represents the PRNU fingerprint, $\mathbf{P}$ denotes the NUAs or patterns with periodic correlations (see Fig. 3) introduced by the in-camera ISP pipeline, and Φ accounts for various additive noise sources. The additive model is motivated by the fact that $\mathbf{P}$ is consistent with dithering-like processes in the ISP (see Sect. 3), which are typically introduced as signal-independent perturbations. While this is a good starting approximation, we acknowledge that potential non-linear interactions between $\mathbf{K}$ and $\mathbf{P}$ are not captured and are left for future work. Under this model, we assume that $\mathbf{K}$, $\mathbf{P}$, and Φ are mutually uncorrelated processes.

Due to the statistical whiteness of the PRNU $\mathbf{K}$ (cf. [18, Sect. II.C]), and under the assumption that the different noise components Φ are also white, the correlation peaks observed in Fig. 3 can be attributed to the embedded pattern $\mathbf{P}$. From the autocorrelation maps shown in Fig. 3, we can infer the block dimensions $B_1 \times B_2$ involved in the block-based processing. For example, for Pattern 1 we determined a block size of 65×60 (i.e., $B_1 = 65$ and $B_2 = 60$).

To analyze the local correlation structure of these 65×60 patches across the frame, Fig. 5 presents the NCC between the top-left 65×60 block of the estimate $\hat{\mathbf{F}}$ obtained from an S10 device containing Pattern 1 (SD06 in Tab. 2) and the remaining samples of $\hat{\mathbf{F}}$. For clarity, the peaks with the largest magnitude are highlighted using white colored markers. As can be observed, these peaks are arranged along parallel diagonal lines. Along each diagonal, every 65×60 patch is correlated with its diagonally adjacent counterpart. The horizontal and vertical separations between peaks in adjacent diagonals are 1023 pixels and 16 pixels, respectively. In total, six diagonal lines are observed, encompassing 181 correlated blocks. The main diagonal is highlighted with a distinct circled marker, as this set of blocks will later prove to be particularly relevant. Repeating the same analysis with any other non-overlapping patch in the image yields the same structure, shifted according to the selected reference block. Although negative correlations are also present (see Fig. 3(a)), their analysis is left for future work.

Based on the observed block-based correlations, and given that we were unable to identify the exact dependency structure among the correlated blocks (despite multiple attempts to reverse-engineer the underlying periodic relationships), we resort to the SVD of the correlated blocks. The objective is to determine whether the *signal subspace* (induced by the block-based dependencies modeled by $\mathbf{P}$) can be separated from the *noise subspace*, which comprises the PRNU component $\mathbf{K}$ and the remaining noise sources Φ .

To this end, we consider a total of N distinct correlated blocks as depicted in Fig. 5 (i.e., $N \leq 181$ in our example). Each block is vectorized and stacked column-wise to form an observation matrix $\mathbf{Z}_N \in \mathbb{R}^{B_1 \cdot B_2 \times N}$. Its SVD can be expressed as

$$\mathbf{Z}_N = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (4)$$

where $\mathbf{U} \in \mathbb{R}^{B_1 B_2 \times B_1 B_2}$ is a unitary matrix whose columns are the left singular vectors of $\mathbf{Z}_N$; $\mathbf{\Sigma} \in \mathbb{R}^{B_1 B_2 \times N}$ is a diagonal (rectangular) matrix whose diagonal entries σ_i , with $i \in \{0, \dots, \min(B_1 \cdot B_2, N) - 1\}$, are the singular values of $\mathbf{Z}_N$, sorted in descending order; and $\mathbf{V} \in \mathbb{R}^{N \times N}$ is a unitary matrix containing the right singular vectors of $\mathbf{Z}_N$. Under this model, the dominant singular vectors (i.e., those associated with the largest singular values) are expected to span the signal subspace induced by the block-based processing, whereas the remaining components span the noise subspace, which includes the PRNU signal $\mathbf{K}$ and the residual noise Φ .

Fig. 6(a) illustrates the evolution of the singular values when the observation matrix $\mathbf{Z}_N$ is constructed using the $N = 181$ blocks corresponding to the six diagonal correlation lines marked with white symbols in Fig. 5. When the correlated blocks are grouped in this manner, the resulting SVD does not exhibit a clear spectral gap that would indicate a transition between the Pattern 1 subspace and the subspace associated with the PRNU and the remaining noise components. In contrast, when $\mathbf{Z}_N$ is formed using only the $N = 46$ blocks associated with the white circled correlation points in Fig. 5 (i.e., those on the main diagonal), the scree plot in Fig. 6(b) exhibits a pronounced gap after the singular value at index $i = 8$. This suggests a signal subspace of rank $r = 9$, with the remaining $N - r = 37$ singular values spanning the PRNU and residual noise subspace.

We attribute the presence or absence of the spectral gap to differences in the strength and coherence of the correlations among the selected blocks. For instance, blocks within the same diagonal exhibit high structural consistency, allowing the NUAs to be compactly represented in a low-dimensional subspace. However, when blocks from different diagonals are included, their relative shifts (although still statistically related) introduce additional structural variability that blurs the separation between signal and noise components, making the spectral gap less pronounced.

Accordingly, we restrict the construction of $\mathbf{Z}_N$ to correlated blocks within the same diagonal (such as the $N = 46$ blocks associated with the white circled correlation points in Fig. 5 for Pattern 1). Then, PRNU recovery is performed by truncating the identified NUAs (signal) subspace, i.e., by setting the first r singular values to zero (for Pattern 1, $r = 9$), such that $\sigma_i = 0$ for $i = 0, \dots, r - 1$, yielding a modified diagonal matrix $\mathbf{\Sigma}'$. Using (4), we compute $\mathbf{Z}'_N = \mathbf{U}\mathbf{\Sigma}'\mathbf{V}^T$, and the estimated PRNU $\hat{\mathbf{K}}'$ is obtained by re-mapping each column of $\mathbf{Z}'_N$ to its original spatial coordinates in $\hat{\mathbf{K}}'$. This procedure is repeated for successive non-overlapping $B_1 \times B_2$ blocks




Figure 6: SVD-based subspace analysis for Pattern 1 from S10 (SD06). Singular values of Z_N using $N = 181$ blocks (no clear spectral gap) in (a). Singular values using $N = 46$ main-diagonal blocks (circled markers in Fig. 5), showing a gap after $i = 8$ and indicating $r = 9$ in (b). Evolution of $\rho(\hat{K}, \hat{K}_r')$ versus r , with a maximum at $r = 9$ and the reference NCC value $\rho(\hat{K}, \hat{F})$ in (c).

along both spatial dimensions in order to cover the entire frame. It is worth noting that, near the image borders, the number of available correlated blocks N used to build Z_N may be smaller than r . In such cases, the PRNU signal cannot be reliably estimated in those regions (in particular, near the top-right and bottom-left corners). Since the number of unrecoverable PRNU samples is expected to represent only a small fraction of the entire frame (e.g., less than 10% of the total image resolution in the case of Pattern 1 with $r = 9$), the corresponding entries of the reconstructed PRNU $\hat{K}'$ are set to zero. As shown later in Sect. 5, this has no significant impact on full-frame source camera verification using the reconstructed PRNU $\hat{K}'$.

Finally, for the Pattern 1 case of the S10 device (SD06), the availability of a raw-based PRNU estimate $\hat{K}$ allows us to validate the choice $r = 9$ (for $N = 46$) by computing the NCC between $\hat{K}$ and the reconstructed PRNU $\hat{K}_r'$ obtained for a given r . Specifically, Fig. 6(c) shows the evolution of $\rho(\hat{K}, \hat{K}_r')$ as a function of r . The maximum NCC is achieved at $r = 9$, as expected. For comparison, the NCC computed without subspace suppression (i.e., between $\hat{K}$ and $\hat{F}$) is also reported. This baseline value provides a useful reference, indicating that the maximum NCC achieved at $r = 9$ is substantially higher and, as further discussed in Sect. 5, is sufficiently high to ensure satisfactory attribution performance.

4.2 PRNU Synchronization under HDR Processing

Although NUA-affected reference estimates $\hat{F}$ give rise to the fingerprint collisions described in Sect. 3, these periodic NUA structures can be exploited to detect local translations introduced during HDR image construction [11, Sect. 2], as the diagonal correlation pattern shifts coherently with the applied transformation.

To illustrate this behavior, we compute the cross-correlation between blocks of the SD06 reference estimate $\hat{F}$ (containing Pattern 1 artifacts) and the residues of both non-HDR and HDR images captured with device SD01 (see examples in Figs. 7(a)–(b)). For the 512×512 top-left block of the non-HDR image (red region in Fig. 7(a)), the correlation map (Fig. 7(c)) shows the dominant peak




Figure 7: Examples of cross-correlation between the estimate $\hat{F}$ (with Pattern 1 NUAs) and image residuals. Top row: non-HDR (a) and HDR (b) images from S9 device [6, C14] (CC BY-SA 4.0). Bottom row: non-HDR top-left block exhibiting a peak at $(0, 0)$ (c); HDR top-left block with shifted peaks and maximum at $(5, 3)$ (d); HDR bottom-left block with maximum at $(5, -3)$ (e), revealing local translations induced by HDR processing.

at $(0, 0)$. In contrast, the corresponding block in the HDR image (Fig. 7(b)) produces multiple nearby peaks of similar magnitude (Fig. 7(d)), with the maximum shifted to $(5, 3)$. For the bottom-left HDR block (Fig. 7(e)), the maximum shifts to $(5, -3)$, revealing a local translation in that region. These observations indicate that a block-wise analysis across the frame can provide a map of local translations.

Based on these findings, we propose a strategy that jointly exploits the NUA-affected estimate $\hat{F}$ and the SVD-reconstructed PRNU $\hat{K}'$ to mitigate misdetections arising from HDR processing.




Figure 8: Block-wise translations $(t_1, t_2)_b$ for an HDR image from a Samsung Galaxy S10 (SD05) device in the FloreView dataset [8, D44] (CC BY-SA 4.0), visualized as a quiver plot overlaid on the block-wise (21×21) NCC map between $\hat{\mathbf{F}}$ and $\mathbf{W}$. The resulting field indicates a global downscaling with coherent shifts. Red dots denote blocks with PCE below τ_b , i.e., $(t_1, t_2)_b = (0, 0)$.

Since $\hat{\mathbf{F}}$ enables local translation estimation, we first compute the residue $\mathbf{W}$ of the image under analysis (HDR or not), divide it into non-overlapping $B \times B$ blocks, and evaluate the PCE as in (2) using the corresponding block of $\hat{\mathbf{F}}$ instead of $\hat{\mathbf{K}}$. As in [11], only blocks whose PCE exceeds a threshold are analyzed and corrected. The per-block translation parameters $(t_1, t_2)_b$ are estimated as

$$(t_1, t_2)_b = \begin{cases} \arg \max_{(s_1, s_2) \in \mathcal{S}_b} \rho_{(s_1, s_2)}(\mathbf{W}_b, \hat{\mathbf{F}}_b \circ \mathbf{Y}_b), & \text{if } \text{PCE}(\mathbf{W}_b, \hat{\mathbf{F}}_b \circ \mathbf{Y}_b) \geq \tau_b \\ (0, 0), & \text{otherwise,} \end{cases}$$

where b indexes the blocks, $\mathcal{S}_b$ denotes the admissible shifts, and τ_b is the decision threshold.

Fig. 8 depicts the estimated local translations $(t_1, t_2)_b$, visualized as a quiver plot, for an image from the FloreView dataset [8] captured by a Galaxy S10 (SD05 in Tab. 2). The block-wise (21×21) NCC map between $\hat{\mathbf{F}}$ and $\mathbf{W}$, smoothed with a 5×5 box filter, is overlaid to highlight regions of strong correlation. The resulting vector field indicates a global downscaling concentrated around these high-correlation areas. Red dots indicate blocks with PCE below the threshold τ_b , for which $(t_1, t_2)_b = (0, 0)$. These typically correspond to highly textured regions, where estimation is less reliable due to weaker pattern presence and noisier residuals. Consequently, in HDR images dominated by strong texture, the translation map becomes less informative. Occasional outliers (e.g., the estimate (56, 63) in Fig. 8) deviate from the general trend; although their impact is limited, refined outlier rejection is left for future work.

Once the translation map is computed, each $B \times B$ patch of the SVD-reconstructed PRNU fingerprint $\hat{\mathbf{K}}'$ is shifted according to its associated $(t_1, t_2)_b$, yielding an HDR-synchronized fingerprint,

denoted as $\hat{\mathbf{K}}'_{\text{HDR}}$. Using $\hat{\mathbf{K}}'_{\text{HDR}}$ in place of $\hat{\mathbf{K}}$, the global PCE is computed as in (2), and the final decision is taken according to (3).

5 Experimental Results

We conducted a series of experiments to validate the proposed framework for source camera verification based on SVD-based PRNU recovery and HDR synchronization. The dataset used both to substantiate the analysis of periodic NUAs in Samsung devices (cf. Sect. 3) and to evaluate the proposed approach is summarized in Tab. 2. It comprises a total of 1,512 images collected from 39 devices across 29 distinct models, all equipped with Exynos chipsets. The table details the number of images per sensor (wide and/or telephoto), distinguishing those employed for fingerprint estimation (ref.) from those reserved for testing (test).

Baseline PRNU-based source camera verification was performed using the Python implementation of the Camera Fingerprint Program [1]. The additional code developed to implement the SVD-based PRNU recovery and HDR synchronization is available at <https://github.com/dvazquezpadin/samsung-NUAs>.

5.1 Source Camera Verification with SVD-Based PRNU Recovery and HDR Synchronization

Building upon the dataset analysis summarized in Tab. 2, six distinct Samsung NUAs were identified in the form of periodic patterns (see Sect. 3). In this subsection, we concentrate on Pattern 1, which is the most prevalent across devices, and experimentally evaluate the proposed SVD-based PRNU recovery (Sect. 4.1) together with the HDR-aware synchronization strategy introduced in Sect. 4.2. Results for one of the remaining patterns are reported in Sect. 5.2.

Here, the evaluation focuses on Galaxy S series devices exhibiting Pattern 1 for which flatfield images captured in default photo mode are available, enabling the calculation of the reference estimate $\hat{\mathbf{F}}$. Specifically, this includes devices SD01, SD03, SD05, SD06, and SD08, corresponding to models S9, S9+, and S10+, including two independent S10 units and the telephoto sensor of one S10 device. For testing, we further incorporate images from additional Pattern 1 devices listed in Tab. 2 and summarized in Tab. 1, including variants from the Note9/Note10 and S20/S21 series.

For each device, the NUA-affected reference estimate $\hat{\mathbf{F}}$ is first estimated and subsequently processed using the SVD-based decomposition described in Sect. 4.1, with block size 65×60 ($B_1 = 65$, $B_2 = 60$), $N = 46$ (main diagonal blocks), and rank truncation $r = 9$, yielding the reconstructed PRNU $\hat{\mathbf{K}}'$. HDR synchronization is then configured using non-overlapping blocks of size $B \times B$ with $B = 256$ and a PCE threshold $\tau_b = 30$, selected empirically and close to the one used in [11]. Performance is evaluated in terms of the Area Under the ROC Curve (AUC) and both the True Positive Rate (TPR) and False Positive Rate (FPR) obtained using the conventional PCE threshold $\tau = 60$. Three configurations are compared: (i) the baseline using the NUA-affected estimate $\hat{\mathbf{F}}$, (ii) the recovered PRNU using the SVD $\hat{\mathbf{K}}'$, and (iii) the HDR-synchronized PRNU $\hat{\mathbf{K}}'_{\text{HDR}}$.

The obtained ROC curves are shown in Fig. 9. The SVD-based recovery markedly improves performance over the baseline, which in some cases (e.g., SD05 and SD08 from FloreView) approaches random guessing (see Figs. 9(c) and 9(f)). However, residual errors remain for HDR images when synchronization is not applied. After




Figure 9: ROC curves for Pattern 1 devices comparing the NUA-affected baseline $\hat{F}$, the SVD-recovered PRNU $\hat{K}'$, and the HDR-synchronized PRNU $\hat{K}'_{HDR}$. SVD recovery outperforms baseline and HDR synchronization yields near-perfect detection.

HDR synchronization, near-perfect detection is achieved for all analyzed sensors. No difference between $\hat{K}'$ and $\hat{K}'_{HDR}$ is observed for device SD03 (S9+) in Fig. 9(b), since no HDR images are available for this sensor. It is also worth noting that the performance achieved with the SVD-recovered PRNU remains remarkably strong for devices SD03 (S9+) and SD01 (S9), even though only 5 flatfield images were used to estimate the initial reference $\hat{F}$ (whereas at least 20 reference images are available for the remaining devices).

Failure cases after synchronization are primarily associated with images dominated by strong texture (e.g., certain samples from SD06 (S10) and SD08 (S10+)), which yield unreliable translation maps. Additionally, one image from SD06 (S10) labeled as MFP3 instead of HDR suggests a different multi-frame fusion process not fully compensated by the proposed method.

Finally, to assess practical performance at the standard threshold $\tau = 60$, we compute the TPR-FPR performance matrices across devices sharing Pattern 1 (only those from which we have flatfield images to extract the fingerprint). Results are shown in Fig. 10(a) for $\hat{K}'$ and in Fig. 10(b) for $\hat{K}'_{HDR}$. The SVD-based recovery effectively

controls false positives, as the FPR remains nearly unchanged when using either $\hat{K}'$ or $\hat{K}'_{HDR}$, whereas HDR synchronization markedly improves the TPR, bringing it close to unity in most cases. As expected, no TPR variation is observed for device SD03 (S9+) due to the absence of HDR images.

5.2 Analysis of other NUA patterns

The proposed SVD-based PRNU recovery framework was subsequently evaluated on the remaining five patterns by applying the same subspace decomposition procedure to their corresponding NUA-affected reference estimates $\hat{F}$. In each case, the block size $B_1 \times B_2$ was adapted to the corresponding periodic structure, and the observation matrix Z_N was formed from the correlated blocks along the main diagonal. However, the results were significantly less satisfactory than those reported in Sect. 5.1 for Pattern 1. In most cases, no clear improvement over the baseline performance obtained with $\hat{F}$ was observed, indicating limited effectiveness of the proposed approach for these patterns. Given the diversity and the increased complexity of the remaining patterns, we concentrate here




Figure 10: TPR-FPR performance matrices at the standard threshold $\tau = 60$ for Pattern 1 devices with available flat-field reference images. (a) SVD-recovered PRNU $\hat{K}'$. (b) HDR-synchronized PRNU $\hat{K}'_{\text{HDR}}$. SVD recovery maintains similar FPR levels, while HDR synchronization substantially increases TPR, approaching unity in most cases.

on Pattern 6[†], which is predominant in the most recent Samsung devices.

In this context, we focused the evaluation on Galaxy S24 devices (SD15, SD16, and SD17 in Tab. 2). The reference estimate $\hat{F}$ was computed according to (1) using $L = 20$ flat-field images captured in default photo mode with the S24 (SD15) wide sensor. For testing, we incorporate all test images from the three devices, having an evaluation set of 20 default photo mode images captured with the same SD15 wide sensor (positives) and 49 default photo mode images captured with SD16 and SD17 wide sensors (negatives).

After estimating the NUA-affected reference pattern $\hat{F}$, we applied the SVD-based decomposition described in Sect. 4.1. However, unlike Pattern 1, devices exhibiting Pattern 6[†] do not display a strictly regular diagonal arrangement of correlated blocks. For Pattern 6 (native resolution 3080×4060), the periodic structure can be described using blocks of size 1×255 ($B_1 = 1$, $B_2 = 255$): correlation peaks along the same diagonal are separated horizontally by B_2 and vertically by 16 pixels (yielding 16 blocks per diagonal), while adjacent diagonals are vertically separated by 257 pixels with no horizontal shift.

Now, for Pattern 6[†] (resolution 3000×4000), the effective block size becomes 1×250 ($B_1 = 1$, $B_2 = 250$). Although horizontal separations remain equal to B_2 , the vertical spacing within a diagonal is not perfectly constant (in some cases 15 instead of 16 pixels), and adjacent diagonals are separated by 252 pixels vertically with no horizontal offset. These deviations indicate that Pattern 6[†] originates from the same underlying structure as Pattern 6, but the image is subsequently resized, altering the geometric regularity of the embedded periodic artifacts. To verify this hypothesis, we resized the reference estimate $\hat{F}$ from SD15 to match the native resolution associated with Pattern 6. Using the corresponding block size 1×255 ($B_1 = 1$, $B_2 = 255$), the expected periodic structure of Pattern 6 was recovered. We then constructed the observation




Figure 11: ROC curves for S24 (SD15) wide sensor evaluation. (a) Default-mode estimate $\hat{F}$ (containing Pattern 6[†]) vs. SVD-reconstructed PRNU $\hat{K}'$, showing no improvement. (b) Raw-based PRNU with resizing and with affine registration; registration achieves perfect detection, confirming PRNU persistence in images captured in default photo mode.

matrix using the $N = 16$ blocks from the main diagonal and applied rank truncation with $r = 1$, corresponding to the most prominent spectral gap, to obtain the reconstructed PRNU $\hat{K}'$.

The resulting ROC curve is shown in Fig. 11(a), together with the curve obtained using the default-mode reference estimate $\hat{F}$. The poor performance clearly indicates that, unlike Pattern 1, the SVD-based approach does not generalize effectively to this pattern. Even after testing different values of r and alternative selections of diagonal blocks, no performance improvement over the baseline was achieved.

At this stage, we investigated whether the PRNU was still present in images captured in default photo mode. Following the procedure described in Sect. 4, we used the PRO mode to acquire raw images

Table 2: Dataset of Exynos-equipped devices used for NUA analysis and experimental validation. The table lists the device ID, model and version, sensor type (wide or telephoto), number of images used for reference pattern estimation (ref.) and testing (test), image resolution ($H \times W$), and data source.

ID	Model	Version	Sensor	#images		Resolution	Source
				ref.	test	(H × W)	
USA Version (Galaxy S Series)							
—	S9	SM-G960U	wide	—	245	3024×4032	[14]
—	S9+	SM-G965U	wide	—	257	3024×4032	[14]
—	S10	SM-G973U	wide	—	133	3024×4032	[14]
—	S10+	SM-G975U	wide	—	228	3024×4032	[14]
Global Version (Galaxy S Series)							
SD01	S9	SM-G960F	wide	5	10	3024×4032	C14 [6]
SD02	S9	SM-G960F	wide	—	9	3024×4032	[4]
SD03	S9+	SM-G965F	wide	5	10	3024×4032	C15 [6]
SD04	S9+	SM-G965F	wide	—	8	3024×4032	[4]
SD05	S10	SM-G973F	wide	54	108	3024×4032	D44 [8]
SD06	S10	SM-G973F	wide	40 [†]	71	3024×4032	Ours
SD07	S10	SM-G973F	telephoto	20	41	3024×4032	Ours
			wide	—	10	3024×4032	[4]
			telephoto	—	6	3024×4032	[4]
SD08	S10+	SM-G975F	wide	32	108	3024×4032	D30 [8]
SD09	S20	SM-G980F	wide	—	11	3024×4032	[4]
SD10	S20+	SM-G985F	wide	—	8	3024×4032	[4]
SD11	S20 Ultra	SM-G988B	wide	—	18	3000×4000	[4]
SD12	S21 5G	SM-G991B	wide	—	6	3000×4000	[4]
SD13	S21+ 5G	SM-G996B	wide	—	8	3000×4000	[4]
SD14	S22 Ultra 5G	SM-S908B	wide	—	12	3000×4000	[4]
SD15	S24	SM-S921B	wide	40 [†]	20	3000×4000	Ours
SD16	S24	SM-S921B	wide	21	37	3000×4000	Ours
SD17	S24	SM-S921B	wide	—	12	3000×4000	[4]
SD18	S24+	SM-S926B	wide	—	16	3000×4000	[4]
SD19	S24 FE	SM-S721B	wide	—	15	3000×4000	[4]
SD20	S25 FE	SM-S731B	wide	—	18	3000×4000	[4]
Chinese Market Version (vivo X60 Series)							
VD01	vivo X60	China	wide	170	82	3000×4000	[12]
VD02	vivo X60 Pro	China	wide	85	81	3000×4000	[12]
Galaxy Note Series							
ND01	Note9	SM-N960F	wide	—	9	3024×4032	[4]
ND02	Note10	SM-N970F	wide	—	12	3024×4032	[4]
ND03	Note10+	SM-N975F	wide	—	12	3024×4032	[4]
ND04	Note10 Lite	SM-N770F	wide	—	6	3024×4032	[4]
Galaxy A Series							
AD01	A25 5G	SM-A256B	wide	—	10	3060×4080	[4]
AD02	A26 5G	SM-A266B	wide	—	10	3060×4080	[4]
AD03	A33 5G	SM-A336B	wide	—	15	3000×4000	[4]
AD04	A35 5G	SM-A356B	wide	—	9	3060×4080	[4]
AD05	A51 5G	SM-A516B	wide	—	12	3000×4000	[4]
AD06	A53 5G	SM-A536B	wide	12	14	3468×4624	[3]
AD07	A53 5G	SM-A536B	wide	—	9	3468×4624	[4]
AD08	A54 5G	SM-A546B	wide	9	17	3060×4080	Ours
AD09	A54 5G	SM-A546B	wide	—	18	3060×4080	[4]
AD10	A55 5G	SM-A556B	wide	22	19	3060×4080	[3]
AD11	A55 5G	SM-A556B	wide	—	25	3060×4080	[4]
AD12	A56 5G	SM-A566B	wide	20	62	3000×4000	Ours
AD13	A56 5G	SM-A566B	wide	—	13	3000×4000	[4]

[†] 20 (out of 40) are captured under the PRO mode.

and estimate the PRNU. Specifically, with device SD15 (S24 wide sensor), we captured $L = 20$ flat-field raw images. Unlike the default photo mode resolution of 3000×4000 , the raw images have a resolution of 3080×4060 , which coincides with that associated with Pattern 6. From these images, we estimated $\hat{\mathbf{K}}$ and computed the PCE as in (2), after resizing $\hat{\mathbf{K}}$ to match the size of the test residue $\mathbf{W}$. The corresponding ROC curve, labeled “w/ resizing,” is

shown in Fig. 11(b). Although this approach yields better results than the SVD-reconstructed PRNU, indicating that traces of the sensor PRNU persist in default photo mode images, the performance remains markedly inferior to that achieved with raw images from the S10 device (cf. Fig. 4).

Since PRO mode provides both JPEG images and their corresponding raw (DNG) counterparts, we observed that simple resizing of the fingerprint is insufficient to properly align the scene content across both versions and, consequently, to synchronize the underlying PRNU. To address this issue, we performed image registration by estimating the affine transformation between the raw and JPEG versions using SURF descriptors. The estimated mapping revealed anisotropic scaling (different horizontal and vertical scale factors) together with a non-zero translation component. Computing the PCE with the geometrically corrected estimate produces the ROC curve labeled “w/ registration” in Fig. 11(b), achieving perfect detection performance. These results finally confirm that the PRNU remains present in default photo mode images of the S24 series.

Since the sensor PRNU remains detectable in default photo mode images, we believe it is feasible to mitigate the impact of this new type of NUAs. However, developing an effective suppression strategy requires a more thorough analysis of the structural dependencies among the blocks associated with Pattern 6[†]. Extending this investigation to Pattern 6[†] and the remaining patterns is left for future work.

5.3 Analysis of Portrait/Bokeh Mode Images

In the initial work by Albisani *et al.* [6], it was suggested that bokeh mode may introduce unexpected correlations among images captured by different Samsung devices. Our analysis, however, indicates that the observed NUAs originate from the default photo mode, as discussed in Sect. 3, rather than from the bokeh effect itself. Unlike Apple devices [22], where fixed synthetic patterns are introduced specifically in blurred regions, Samsung portrait images do not exhibit correlation between bokeh areas and the reference estimate $\hat{\mathbf{F}}$, nor between different images. To verify this, we computed a block-wise NCC map (block size 21×21) between the default-mode estimate $\hat{\mathbf{F}}$ (containing Pattern 1 NUAs) and the residual of a portrait image (illustrated in Fig. 12(a)). The resulting NCC map (Fig. 12(b)) shows no significant correlation in bokeh regions, whereas the corresponding original (pre-bokeh) image (Fig. 12(c)) displays a relatively uniform NCC map across the frame (Fig. 12(d)). Although not reported here for brevity, we also verified that bokeh regions do not exhibit mutual correlation across images.

These observations confirm that bokeh processing does not introduce additional false positives. However, it may still degrade detection performance when using a valid PRNU estimate, since the low-pass filtering and noise injection employed to generate the bokeh effect attenuate the PRNU component in the residue $\mathbf{W}$. Therefore, excluding bokeh regions during both reference estimation and verification is advisable, especially when only portrait images are available. Investigating strategies that exploit knowledge of the NUA patterns to compensate for bokeh regions is left for future work.




Figure 12: Example of a block-wise NCC map (21×21 blocks) between the estimate $\hat{F}$ (including Pattern 1 NUAs) and image residuals from bokeh and pre-bokeh images. (a) Portrait (bokeh) image from [6, C14] (CC BY-SA 4.0). (b) Corresponding NCC map showing no correlation in bokeh-affected regions. (c) Original pre-bokeh image from [6, C14] (CC BY-SA 4.0). (d) Corresponding NCC map exhibiting uniform correlation across the frame.

6 Conclusions

This work investigated NUAs in images captured by Samsung smartphones equipped with Exynos and their impact on PRNU-based source camera verification. Across 39 devices and 29 models (released between 2018 and 2025), we identified 6 distinct diagonal periodic patterns that induce fingerprint collisions and significantly degrade attribution performance. Our study clarifies that these patterns stem from the default photo mode ISP processing, not from the sensor itself, which is shared with Qualcomm-based USA variants exhibiting different NUA patterns.

Our key contribution is the introduction of a subspace-based suppression strategy that leverages SVD to isolate and remove NUAs while preserving the underlying sensor PRNU, thereby restoring reliable PRNU-based attribution. For the most prevalent case (Pattern 1), the proposed approach yields a substantial improvement over the baseline. We hope that this formulation will stimulate further research on the increasingly ubiquitous NUAs introduced by modern smartphone imaging pipelines. Moreover, as a by-product of the regular structure of these NUAs, we exploit the artifact itself to address HDR-induced misdetections. Using the periodic pattern to estimate local geometric shifts and synchronize the reconstructed PRNU, we effectively compensate for HDR-related misalignments, achieving near-perfect detection performance.

The proposed approach, however, does not generalize to all identified patterns. In particular, for more recent structures such as Pattern 6[†], the increased structural complexity and the geometric distortions introduced by in-camera resizing hinder effective subspace separation. Nevertheless, experiments conducted with raw images confirm that the underlying PRNU sensor remains embedded in default photo mode acquisitions. This evidence suggests that effective mitigation should be achievable and motivates future research toward a unified framework capable of addressing all observed NUA patterns.

Acknowledgments

This work is supported in part by Xunta de Galicia and the European Regional Development Fund under Project ED431C 2025/41.

References

- [1] [n. d.]. Camera Fingerprint - Matlab & Python implementation. https://dde.binghamton.edu/download/camera_fingerprint/
- [2] [n. d.]. dcrow. <https://dechifro.org/dcrow>
- [3] [n. d.]. Flickr API. <https://www.flickr.com/services/api/>
- [4] [n. d.]. GSMArena. <https://www.gsmarena.com/>
- [5] [n. d.]. Samsung's 108Mp ISOCELL Bright HM1 delivers brighter ultra-high-res images with industry-first Nonapixel technology. <https://news.samsung.com/global/samsungs-108mp-isocell-bright-hm1-delivers-brighter-ultra-high-res-images-with-industry-first-nonapixel-technology>
- [6] C. Albisani, M. Iuliani, and A. Piva. 2021. Checking PRNU Usability on Modern Devices. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 2535–2539. doi:10.1109/ICASSP39728.2021.9413611
- [7] Daniele Baracchi, Massimo Iuliani, Andrea G. Nencini, and Alessandro Piva. 2021. Facing Image Source Attribution on iPhone X. In *Digital Forensics and Watermarking*. Xianfeng Zhao, Yun-Qing Shi, Alessandro Piva, and Hyoung Joong Kim (Eds.). Springer International Publishing, Cham, 196–207.
- [8] Daniele Baracchi, Dasara Shullani, Massimo Iuliani, and Alessandro Piva. 2023. FloreView: An Image and Video Dataset for Forensic Analysis. *IEEE Access* 11 (2023), 109267–109282. doi:10.1109/ACCESS.2023.3321991
- [9] Jan Butora and Patrick Bas. 2024. The Adobe Hidden Feature and its Impact on Sensor Attribution. In *Proceedings of the 2024 ACM Workshop on Information Hiding and Multimedia Security* (Baiona, Spain). Association for Computing Machinery, New York, NY, USA, 143–148. doi:10.1145/3658664.3659650
- [10] Mo Chen, Jessica Fridrich, Miroslav Goljan, and Jan Lukáš. 2008. Determining Image Origin and Integrity Using Sensor Noise. *IEEE Transactions on Information Forensics and Security* 3, 1 (2008), 74–90. doi:10.1109/TIFS.2007.916285
- [11] Morteza Darvish Morshedi Hosseini and Miroslav Goljan. 2019. Camera Identification from HDR Images. In *Proceedings of the ACM Workshop on Information Hiding and Multimedia Security* (Paris, France) (IH&MMSec'19). Association for Computing Machinery, New York, NY, USA, 69–76. doi:10.1145/3335203.3335717
- [12] Suwen Du, Pengpeng Yang, Daniele Baracchi, Jinglian Jin, Dasara Shullani, and Alessandro Piva. 2025. ForensicCam-215K: A Large Scale Image and Video Dataset for Forensic Analysis. In *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 1–5. doi:10.1109/ICASSP49660.2025.10890764
- [13] Miroslav Goljan, Jessica Fridrich, and Tomáš Filler. 2009. Large scale test of sensor fingerprint camera identification. In *Media Forensics and Security*, Edward J. Delp III, Jana Dittmann, Nasir D. Memon, and Ping Wah Wong (Eds.), Vol. 7254. International Society for Optics and Photonics, SPIE, 72540L. doi:10.1117/12.805701
- [14] Massimo Iuliani, Marco Fontani, and Alessandro Piva. 2021. A Leak in PRNU Based Source Identification. Questioning Fingerprint Uniqueness. *IEEE Access* 9 (2021), 52455–52463. doi:10.1109/ACCESS.2021.3070478
- [15] M. Kivanc Mihcak, I. Kozintsev, and K. Ramchandran. 1999. Spatially adaptive statistical modeling of wavelet image coefficients and its application to denoising. In *1999 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings. ICASSP99 (Cat. No.99CH36258)*, Vol. 6. 3253–3256 vol.6. doi:10.1109/ICASSP.1999.757535
- [16] Samantha Klier and Harald Baier. 2025. Source Camera Identification - Do we have a gold standard? *Forensic Science International: Digital Investigation* 52 (2025), 301858. doi:10.1016/j.fsidi.2024.301858
- [17] Jan Lukáš, Jessica Fridrich, and Miroslav Goljan. 2005. Determining digital image origin using sensor imperfections. In *Image and Video Communications and Processing 2005*, Amir Said and John G. Apostolopoulos (Eds.), Vol. 5685. International Society for Optics and Photonics, SPIE, 249 – 260. doi:10.1117/12.587105
- [18] J. Lukáš, J. Fridrich, and M. Goljan. 2006. Digital camera identification from sensor pattern noise. *IEEE Transactions on Information Forensics and Security* 1, 2 (2006), 205–214. doi:10.1109/TIFS.2006.873602
- [19] Andrea Montibeller, Roy Alia Asiku, Fernando Pérez González, and Giulia Boato. 2024. Shedding Light on some Leaks in PRNU-based Source Attribution. In *Proceedings of the 2024 ACM Workshop on Information Hiding and Multimedia Security* (Baiona, Spain) (IH&MMSec'24). Association for Computing Machinery, New York, NY, USA, 137–142. doi:10.1145/3658664.3659654
- [20] Andrea Montibeller and Fernando Pérez-González. 2024. An Adaptive Method for Camera Attribution Under Complex Radial Distortion Corrections. *IEEE Transactions on Information Forensics and Security* 19 (2024), 385–400. doi:10.1109/TIFS.2023.3318933
- [21] Nabeel Nisar Bhat and Tiziano Bianchi. 2022. Investigating Inconsistencies in PRNU-Based Camera Identification. In *2022 IEEE International Conference on Image Processing (ICIP)*. 851–855. doi:10.1109/ICIP46576.2022.9897923
- [22] David Vázquez-Padín, Fernando Pérez-González, and Pablo Pérez-Miguel. 2026. Apple's Synthetic Defocus Noise Pattern: Characterization and Forensic Applications. *IEEE Transactions on Information Forensics and Security* 21 (2026), 1096–1111. doi:10.1109/TIFS.2026.3653213