

Technical Report: Expanded Analysis and Experimental Evaluation of Apple’s Synthetic Defocus Noise Pattern

David Vázquez-Padín, Fernando Pérez-González, and Pablo Pérez-Miguelé*^{*}

Summary

This technical report complements the work in [1], with expanded details and supporting analyses. Specifically:

- **Section 1** begins by giving further details on the evaluation of the baseline PRNU matching approach applied to the three iPhone 11 Pro users who exhibited PRNU collisions in Iuliani *et al.*’s dataset [2].
- **Section 2** provides complementary insights into Apple’s BP extraction under the NL portrait lighting mode. This includes justifying the denoising parameter K together with the analysis of the correspondence between the BP estimates $\hat{\mathbf{P}}_{\text{NL}}$ and $\hat{\mathbf{P}}_{\text{SLM}}$ (Sect. 2.1), and guidelines for improving NL-based BP extraction (Sect. 2.2).
- **Section 3** expands on the SLM portrait lighting mode, covering the determination of the constant background value (Sect. 3.1), an assessment of BP Gaussianity using a Kolmogorov-Smirnov test (Sect. 3.2), and a look-up table with estimated ISO-dependent scaling factors (Sect. 3.3).
- **Section 4** presents a quality assessment of the BP estimates, including an analysis of Apple’s BP across different resolutions and aspect ratios (Sect. 4.1), additional BP characteristics (Sect. 4.2), and the variability of the brightness-dependent function $g(\cdot)$ across devices (Sect. 4.3).
- **Section 5** expands on the experiments presented in the main paper [1], detailing the threshold selection strategy for BP detection and localization (Sects. 5.1 and 5.2, respectively), the replication of Baracchi *et al.*’s original results (Sect. 5.3), a proof-of-concept for forgery localization using BP-aware PRNU analysis (Sect. 5.4), BP detection “in the wild” using the FFHQ dataset (Sect. 5.5), and additional details on BP detection under post-processing (Sect. 5.6).

For a complete description of the notation used throughout this document, please refer to [1].

^{*}Signal Theory and Communications Department, atlanTTic, Universidade de Vigo, Vigo, Spain (e-mail: dvazquez@gts.uvigo.es; fperez@gts.uvigo.es; pmperez@gts.uvigo.es)




Figure 1: Values of η for the 3 iPhone 11 Pro users with PRNU collisions in [2]. From left to right, each user’s PRNU fingerprint is tested against the other two. The upper pie charts show the distribution of Reference and Test images per user, indicating the proportion of Portrait and Photo images.

1 Baseline PRNU Matching in [2]

As observed for user 102027268@N08 in Figure 1, false positives for users 102054399@N08 and 101543825@N07 also occur exclusively with images captured in Apple’s portrait mode. However, their positive cases exhibit distinct behaviors, which are examined below.

For user 102054399@N08, the similarity measure η in positive cases (consisting solely of photo-mode images) is significantly lower than that of user 102027268@N08, who exclusively captures portrait-mode images. This indicates that while portrait images in the Reference set degrade the PRNU estimate, their absence in the Test set prevents η from achieving a huge magnitude. Furthermore, since 102054399@N08’s Reference set contains only 4 photo-mode images out of 32, the resulting PRNU estimate is affected by traces of Apple’s SDNP, leading to η values close to the threshold.

User 101543825@N07 presents a different scenario. Metadata analysis shows that five photo-mode images were captured using a different LensID (iPhone 11 Pro back triple camera 6mm f/2 vs. iPhone 11 Pro back dual camera 6mm f/2), suggesting the use of a different sensor (see [3]). Additionally, two of these images were taken with digital zoom, as indicated by the EXIF tag DigitalZoomRatio. These 5 images, distributed

Camera	Image ($H \times W$)	Depth Map	Auxiliary Matte
Rear (24MP)	5712×4284	768×576	2016×1512
Rear (12MP)	4032×3024	768×576	2016×1512
Front (7MP)	3088×2316	640×480	1544×1158
[†] Front (7MP)	3088×2320	640×480	1544×1160

[†]Older iPhone models, such as iPhone SE (2nd generation) and iPhone X.

Table 1: Resolution of Images Captured in Apple’s Portrait Mode




Figure 2: NCC values between $\hat{\mathbf{P}}_{\text{SLM}}$ and $\hat{\mathbf{P}}_{\text{NL}}$ for various values of K (a). Tests were conducted for all three available resolutions of the iPhone 15 (24MP, 12MP, and 7MP). Correlation map between $\hat{\mathbf{P}}_{\text{SLM}}$ and $\hat{\mathbf{P}}_{\text{NL}}$ ($K = 33$) for 24MP resolution from the iPhone 15 (b).

across the Reference (2) and Test (3) sets, likely contributed to weaker correlation in positive cases due to spatial misalignment or differences in sensor characteristics, which are not accounted for when extracting or matching the PRNU fingerprint.

2 NL-based BP extraction

To complement Apple’s portrait mode description in [1, Sect. II-B1], Table 1 presents the resolutions of various images captured in portrait mode, including the portrait image, depth map, and auxiliary mattes.

2.1 Correspondence between $\hat{\mathbf{P}}_{\text{NL}}$ and $\hat{\mathbf{P}}_{\text{SLM}}$

As discussed in [1, Sect. III], Apple’s BP can be extracted using two portrait lighting modes: NL and SLM, each with advantages and limitations. The estimate $\hat{\mathbf{P}}_{\text{NL}}$ is less prone to saturation artifacts (unless the background itself is saturated), whereas $\hat{\mathbf{P}}_{\text{SLM}}$ contains such artifacts. However, extracting $\hat{\mathbf{P}}_{\text{SLM}}$ is a simpler process, ensuring no contamination from grain or background texture. Additionally, the denoising step in [1, Eq. (7)] introduces artifacts in $\hat{\mathbf{P}}_{\text{NL}}$ that are absent in $\hat{\mathbf{P}}_{\text{SLM}}$.

To assess the impact of the denoising parameter K on $\hat{\mathbf{P}}_{\text{NL}}$, we compute the NCC, as defined in [1, Eq. (1)], between $\hat{\mathbf{P}}_{\text{SLM}}$ and $\hat{\mathbf{P}}_{\text{NL}}$, i.e., $\rho(\hat{\mathbf{P}}_{\text{SLM}}, \hat{\mathbf{P}}_{\text{NL}})$, for different values of K . Figure 2a shows how the NCC varies




Figure 3: Picture of the setup used to estimate Apple’s BP.

with K , using $L = 8$ across the three available iPhone 15 resolutions. The NCC increases with K , peaking at different values for each resolution ($K = 87$ (24MP), $K = 63$ (12MP), and $K = 51$ (7MP)) before declining. Although an NCC of 1 is unattainable due to the different extraction methods, values above 0.8 are achieved for larger K , indicating strong correlation. Lower values of K are unsuitable for flat backgrounds, as they can introduce excessive image content, reducing the NCC. While the optimal K depends on resolution, we use $K = 33$ for all three, as it provides satisfactory results across the board.

To analyze local variations in correlation between $\hat{\mathbf{P}}_{\text{SLM}}$ and $\hat{\mathbf{P}}_{\text{NL}}$, we generate an NCC map by computing local NCC values over distinct 16×16 blocks across the whole BPs. The resulting map (Figure 2b) shows no significant loss of correlation, with values ranging from 0.75 to nearly 1 and an average around 0.85, as seen in Figure 2a. This suggests that while the two SDNP estimates exhibit distinct characteristics, they still represent the same underlying structure.

2.2 NL-based BP Extraction Guidelines

As noted in [1, Sect. III-A], careful selection of background and capture settings is crucial for BP extraction in the NL scenario to obtain a good estimate $\mathbf{P}$. We recommend the following guidelines:

1. **Consistent ISO:** Keeping a fixed ISO across all 2L images is crucial for maintaining the validity of [1, Eq. (10)], as variations in ISO scale $\mathbf{P}$ differently

and affect quantization noise in [1, Eq. (9)], complicating the model. The impact of such variations is analyzed below in Sect. 4. Extreme ISO values should be avoided to prevent pixel saturation.

2. **Uniform background:** A uniform background is crucial for accurate BP extraction. While painted walls or cardboard are suitable options, they can introduce unwanted texture or grain into the residuals. A more controlled alternative is using a screen displaying a uniform color, as it allows for adjustable luminance. However, Moiré patterns must be avoided. Since $g(y'_l)$ is unlikely to remain perfectly constant across all $2L$ images, it is important to lock the exposure settings (indicated by AE/AF Lock on iPhones) and maintain consistent indoor lighting when capturing both scenes.
3. **Reduced aperture:** Using the minimum aperture (f/1.4), as shown in [1, Fig. 3a], helps minimize surface texture interference and reduces content leakage, enhancing the accuracy of the approximation $H_K(\mathbf{Z}_l) \approx \mathbf{Y}'_l$.
4. **Zoom usage:** Zoom functionality up to 3x in portrait mode (available in recent iPhones such as the iPhone 15) can help minimize background luminance gradients, enhancing the accuracy of the approximation $G(\mathbf{Y}'_l) \approx G(y'_l \cdot \mathbf{1})$. Importantly, we confirmed that, unlike PRNU, zooming does not cause BP misalignment.

By incorporating these considerations, we found that an effective setup for reliable BP fingerprint estimation in the NL scenario involves capturing portrait images with a synthetic face (e.g., a doll) as the foreground and a screen displaying a uniform color as the background, as illustrated in Figure 3.

3 SLM-based BP extraction

3.1 Determining the constant background value in the SLM lighting mode

On the iPhone 15, SLM portrait images cannot be directly captured at 24MP resolution using the camera app; instead, when a portrait is taken in SLM lighting mode at 24MP, it is automatically saved at 12MP resolution. However, the lighting effect can be applied afterward through the Photos app while preserving the 24MP resolution, which is the approach we followed. Regardless, the conclusions drawn here concerning the value of the constant background in SLM lighting mode remain consistent across all available resolutions.

We have empirically observed that under low-light conditions, the depth map generation becomes less accurate, which in turn affects the accuracy of the image regions




Figure 4: A portrait image captured under low-light conditions at ISO 4000 (a), along with its corresponding depth map (b), portrait effects matte (c), and BP-driven NCC map from $\hat{\mathbf{P}}_{\text{NL}}$ (d).




Figure 5: SLM portrait image (a) and its histogram-equalized version, with an overlay mask highlighting in red the original pixels with a value of 4 (b). HKLM portrait image (c) and its histogram-equalized version, with an overlay mask highlighting in red the original pixels with a value of 255 (d).

selected for blurring (see [1, Sect. II-B1]). Figure 4 illustrates this effect with an example of an original portrait image captured at ISO 4000 (Figure 4a), along with the corresponding noisy depth map (Figure 4b), the resulting portrait effects matte (Figure 4c), and the NCC map with a specific BP estimate $\hat{\mathbf{P}}_{\text{NL}}$ (Figure 4d). The NCC map clearly reveals segmentation errors in the foreground, as a significant portion of the actual background remains unblurred.

After confirming that the foreground/background segmentation is imperfect in this case, we apply the SLM lighting effect, resulting in the image shown in Figure 5a. Upon equalizing the image histogram, we observe a background region with a constant value where the SDNP has not been added. For better visibility, this region is highlighted in red in Figure 5b. This constant-value region, which is part of the background rather than the subject in focus and lacks any SDNP, likely corresponds to the underlying $\mathbf{Y}'$ in the model in [1, Eq. (6)]. The observed constant value is 4.

It is worth noting that iPhone devices also offer a lighting effect complementary to the SLM mode, called **High-Key Light Mono** (HKLM). Similar to SLM, HKLM produces a monochromatic effect; however, instead of a dark background, it replaces the background with a solid white one (see an example in Figure 5c). In this case, the constant value used (highlighted in red in the equalized

image, Figure 5d) is 255. This extremely high value results in even more saturated pixels than in SLM mode, which is the main reason we have ruled out using HKLM to obtain another estimate of $\mathbf{P}$. However, we leave for future research the possibility of leveraging both modes to obtain a more accurate estimate of the true BP, as at certain ISO levels pixels that saturate in one mode (SLM) are unlikely to saturate in the other mode (HKLM), and vice versa.

3.2 Assessing BP Gaussianity with the Kolmogorov-Smirnov Test

In order to assess whether the samples of $\mathbf{P}$ can be modeled by a Gaussian distribution, we conducted a series of tests using images collected in the SLM mode for different ISO values. We first describe the procedure for a generic ISO value and subsequently report the results for various ISO settings. We begin with a single sample $\mathbf{Z}_{\text{SLM}}$ captured under a specific ISO setting. To eliminate potential interactions between saturation and compression, we discard all pixels with zero values. From the first 10^6 remaining pixels we compute the empirical pmf which is denoted by $\mathbf{h}_{\text{SLM}}$. Next, we estimate the standard deviation of the underlying Gaussian distribution (i.e., in the absence of compression or saturation) by generating i.i.d. samples of the random variable

$$Z = \text{round}(4 + \sigma_Z P), \quad (1)$$

where $P \sim \mathcal{N}(0, 1)$. We then remove all samples taking values below 1 and keep the first $N_Z = 10^6$ remaining ones. The empirical pmf of such samples is denoted by $\mathbf{h}_Z$. Finally, the estimate of the std, $\hat{\sigma}_Z$, is obtained via an exhaustive search that minimizes the chi-square distance between $\mathbf{h}_{\text{SLM}}$ and $\mathbf{h}_Z$.

Table 2 presents, for each considered ISO setting, the values of the Kolmogorov-Smirnov (KS) statistic used to compare the cumulative density functions of $\mathbf{h}_{\text{SLM}}$ and $\mathbf{h}_Z$, where the latter is obtained with $\sigma_Z = \hat{\sigma}_Z$ in (1) for the given ISO value. While the KS-statistic values are very small, they remain large enough to reject the null hypothesis that the non-zero pixels of $\mathbf{Z}_{\text{SLM}}$ follow the same distribution as the positive-valued samples generated according to (1).

For reference, let $\mathbf{h}_Z$ and $\mathbf{h}'_Z$ be the empirical pmfs corresponding to independently generating two sets of i.i.d. samples of the random variable Z as in (1), and retaining for each set the first 10^6 positive-valued samples. The KS-statistics comparing these distributions across all ISO values have a mean of 8.1707×10^{-4} and a standard deviation of 2.1187×10^{-4} . Furthermore, as the sample size N_Z increases, the KS-statistic comparing $\mathbf{h}_Z$ and $\mathbf{h}_{\text{SLM}}$ does not decrease, contrary to what would occur if the two distributions were truly identical (as observed when comparing $\mathbf{h}_Z$ and $\mathbf{h}'_Z$).

ISO	50	64	100	125	200
KS	0.0033	0.0044	0.0044	0.0076	0.0085
ISO	320	400	500	640	800
KS	0.0072	0.0072	0.0084	0.0074	0.0078
ISO	1000	2000	3200	4000	5000
KS	0.0082	0.0074	0.0071	0.0073	0.0075

Table 2: ISO values and KS statistics.




Figure 6: Cumulative density functions of 10^6 positive-valued pixels of $\mathbf{Z}_{\text{SLM}}$ (ISO 500) and 10^6 i.i.d. samples of a random variable obtained by quantizing a $\mathcal{N}(4, \sigma_Z)$ random variable and discarding non-positive values.

These results suggest that while the underlying distribution of the pixels in $\mathbf{P}$ is not strictly Gaussian, it can be adequately modeled as such. Figure 6 represents the cumulative density functions of $\mathbf{h}_Z$ and $\mathbf{h}_{\text{SLM}}$ for ISO 500, the setting that exhibits the maximum distance in Table 2. Even in this case, the match is excellent, reinforcing the validity of the Gaussian approximation.

Given this close agreement but the simultaneous failure to satisfy the KS-test for a large sample size N_Z and a reasonable p-value (e.g., 0.01), one may question the source of this discrepancy. The stability of the KS-statistics comparing $\mathbf{h}_Z$ and $\mathbf{h}_{\text{SLM}}$ as $N_Z \rightarrow \infty$ suggests that the two distributions are very similar yet subtly different. While we do not explore this hypothesis further in this work, we conjecture that the discrepancy may arise from differences in the tails of the distribution, potentially because the BP might not be directly generated from a Gaussian process but rather by filtering a non-Gaussian process (e.g., a Bernoulli process).

3.3 Expanded Estimation of γ_{ISO}

The exhaustive search conducted in [1, Sect. IV-A] not only empirically validates the assumption that the mean

ISO	40	50	64	100	125	200	250	320	400	500	640	800	1000	2000	3200	4000	5000
$\hat{\gamma}_{\text{ISO}}$	1.363	1.395	1.454	1.791	2.006	2.293	2.4791	2.685	2.718	3.182	3.229	3.269	3.479	3.740	3.885	3.891	3.891

Table 3: Look-Up Table for Estimated $\hat{\gamma}_{\text{ISO}}$ Values at Different ISO Settings.




Figure 7: Contour of cost functions derived from comparing the pmfs of synthetic samples with those from the background of an SLM image at ISO 500, using the KLD (a) and chi-square distance (b).

value of the BP $\mathbf{P}$ is approximately zero but also enables the estimation of the ISO-dependent scaling factor γ_{ISO} for different ISO values. This search is performed over a predefined grid of candidate pairs (μ_Z, σ_Z) . For each candidate pair, synthetic data is generated using a Gaussian model, and the resulting distribution is compared to the observed data distribution (i.e., the empirical pmf of samples from the background of SLM images) using the KLD and chi-square distance.

Figure 7 presents the corresponding cost functions obtained for the particular case of ISO 500 for the KLD (Figure 7a) and the chi-square distance (Figure 7b). In both cases, the plots reveal a well-defined, likely convex cost function with a distinct global minimum (highlighted with a red dot). The shape of the cost functions suggests that the proposed exhaustive search formulation is appropriately posed, reinforcing the reliability of the obtained estimates.

Table 3 gathers the estimated $\hat{\gamma}_{\text{ISO}}$ values at different ISO settings.

4 Quality assessment of BP estimates

As discussed in [1, Sect. III], the quality of the BP estimates derived from NL- or SLM-based extraction methods can vary with factors such as scene capture conditions, ISO settings, and number of captured images L . Although increasing L helps reduce noise, the ISO scaling factor γ_{ISO} and the brightness-dependent scaling function $g(\cdot)$ play a crucial role. To quantify the quality of a given BP estimate $\hat{\mathbf{P}}$ (coming either from $\hat{\mathbf{P}}_{\text{NL}}$ or $\hat{\mathbf{P}}_{\text{SLM}}$), we use the normalized cross-correlation with a given residual block $\mathbf{W}$ of size $B \times B$ as specified in [1, Eq. (7)], i.e., $Q(\mathbf{W}) \triangleq \rho(\hat{\mathbf{P}}, \mathbf{W})$. This residual corresponds to a

flat-background block $\mathbf{Y}'$ characterized by its mean value $y' = \mu(\mathbf{Y}')$. In principle, the larger $Q(\mathbf{W})$, the higher the quality of the estimate of $\mathbf{P}$. In the sequel, we will use the subindices NL and SLM on $Q(\mathbf{W})$ to specify which of the two estimation methods is applied.

Let us analyze first the value of $Q_{\text{NL}}(\mathbf{W})$ for the case where B^2 is large. Under the assumption of ergodicity in the noise process, it is possible to resort to the law of large numbers to show that as $B \rightarrow \infty$:

$$Q_{\text{NL}}(\mathbf{W}) \rightarrow \left(1 + \frac{\text{NBR}_{\text{NL}}}{\left(\gamma_{\text{ISO}}^{(\text{NL})} \cdot \bar{\lambda}_{\text{NL}}\right)^2}\right)^{-\frac{1}{2}} \cdot \left(1 + \frac{\text{NBR}(\mathbf{W})}{\gamma_{\text{ISO}}^2 g^2(y')}\right)^{-\frac{1}{2}}, \quad (2)$$

where we have introduced the superscript in $\gamma_{\text{ISO}}^{(\text{NL})}$ to differentiate it from γ_{ISO} (which corresponds to the residual $\mathbf{W}$) and the following Noise-to-BP Ratios (NBR):

1. $\text{NBR}_{\text{NL}} \triangleq \frac{\mathbb{E}[\|\Psi_{\text{NL}}\|_{\text{F}}^2]}{\|\mathbf{P}\|_{\text{F}}^2}$, which is specific to the NL-based extraction method and
2. $\text{NBR}(\mathbf{W}) \triangleq \frac{\mathbb{E}[\|\Psi\|_{\text{F}}^2]}{\|\mathbf{P}\|_{\text{F}}^2}$, which depends on the noise Ψ contained in residual $\mathbf{W}$.

Note that if we assume that $L \rightarrow \infty$, this leads to $\mathbb{E}[\|\Psi_{\text{NL}}\|_{\text{F}}] \rightarrow 0$ and, consequently, $\text{NBR}_{\text{NL}} \rightarrow 0$, leaving only the second term in (2). This means that the quality increases with L , as one would expect. Furthermore, larger values of the scaling factors γ_{ISO} , $g(y')$ and γ_{NL} are expected to enhance quality.

Similar definitions and expressions apply for the SLM-based BP estimate $\hat{\mathbf{P}}_{\text{SLM}}$. However, notice that, in contrast to the NL case, one can no longer claim that $\mathbb{E}[\|\Psi_{\text{SLM}}\|_{\text{F}}] \rightarrow 0$ as $L \rightarrow \infty$ because, following the procedure described in [1, Sect. III-B], all the L images $\mathbf{Z}$ used for the estimation correspond to $\mathbf{Y}' = 4 \cdot \mathbf{1}$ and the same ISO, so their corresponding residuals (and in turn their noise terms) are very correlated for different realizations.

To empirically evaluate the quality of the estimates, we analyzed the $M \cdot T$ blocks extracted from the ColorChecker portrait images described in [1, Sect. IV] (both for ISO 125 and ISO 400). For each residual block $\mathbf{W}_{m,t}$, with $m = 1, \dots, M$ and $t = 1, \dots, T$, and its spatially aligned block $\hat{\mathbf{P}}_m$ of the BP estimate under analysis, we calculated the sample NCC (as defined in [1, Eq. (1)]). The quality measure $\bar{Q}$ was then determined as the average quality value across all processed blocks, given by:

$$\bar{Q} \triangleq \frac{1}{M \cdot T} \sum_{m=1}^M \sum_{t=1}^T \rho(\hat{\mathbf{P}}_m, \mathbf{W}_{m,t}). \quad (3)$$




Figure 8: Quality metric $\bar{Q}$ as a function of L for different BPs extracted under NL and SLM lighting modes and various ISO settings for the observed blocks: ISO 125 (a), ISO 400 (b). The legend in (b) also applies to (a).

This was done using BP estimates obtained across various L values while capturing flat-background scenes with different average luminance levels under the following conditions:

1. Using the NL lighting mode at ISO 125 and 400:
 - $\hat{\mathbf{P}}_{\text{NL}} (\text{Hi})$: High luminance values (around 180).
 - $\hat{\mathbf{P}}_{\text{NL}} (\text{Mid})$: Medium luminance values (around 90).
2. Using the SLM lighting mode at ISO 125 and 400:
 - $\hat{\mathbf{P}}_{\text{SLM}}$.
3. Using the NL lighting mode:
 - $\hat{\mathbf{P}}_{\text{NL}} (\text{Misc})$: Miscellaneous luminance (between 100 and 200) and ISO values (specifically, 125, 200, 250, 320, 400, and 500).

Figure 8 presents the obtained values of $\bar{Q}$ as a function of L for each case. In general, higher ISO levels in the observed blocks correspond to higher quality values, as seen by comparing Figure 8b with Figure 8a. This trend aligns with the influence of γ_{ISO} in (2).

For BPs extracted using the NL lighting mode, increasing L consistently improves quality, as anticipated. When the ISO values of the observed blocks match those used for BP extraction, higher luminance levels lead to greater quality compared to medium luminance levels. This is because the brightness-dependent scaling function $g(\cdot)$, and thus $\bar{\lambda}_{\text{NL}}$, are larger in the former case. While low-luminance captures should theoretically benefit from a larger $g(\cdot)$, they actually show lower quality. These results are omitted from Figure 8 for clarity. This degradation is likely due to systematic errors like color banding and content leakage, which are more common in low-light conditions. These artifacts not only affect SDNP addition but also compromise scene segmentation from the depth map, as discussed in Section 3.1, which unavoidably leads

to a lower quality estimate of the underlying BP. Finally, we also examined the performance of a BP extracted using different luminance and ISO values, i.e., $\hat{\mathbf{P}}_{\text{NL}} (\text{Misc})$. This extracted BP shows strong performance for small values of L and is only outperformed by the BP extracted at high luminance and matching ISO values for larger L . This indicates that sharing capturing settings, such as the ISO value, induces similar noise components $\Psi_{m,t}$ (i.e., the noise present in the residual $\mathbf{W}_{m,t}$) and $\bar{\Psi}_{\text{NL}}$ (i.e., the average noise from the BP estimation). These noise components exhibit a positive correlation when computing $\rho(\hat{\mathbf{P}}_{\text{NL},m}, \mathbf{W}_{m,t})$ in (3), leading to an overestimation of quality $\bar{Q}_{\text{NL}}$.

Unlike the NL case, increasing L for BPs extracted in the SLM mode does not improve quality. This is because each additional image sample used for BP estimation introduces a highly correlated residual with previous observations, preventing NBR_{SLM} from reaching zero, as previously discussed. Additionally, while higher ISO values would theoretically enhance the metric $\bar{Q}_{\text{SLM}}$, Figure 8 shows that this is not the case, as worse performance is always obtained for ISO 400 with respect to ISO 125. The reason, as illustrated in [1, Fig. 4], is that higher ISO levels lead to more saturated pixels, which amplify the noise component $\mathbb{E}[\|\bar{\Psi}_{\text{SLM}}\|_F]$. Although the signal component also increases with higher ISO, its increase is smaller than that of the noise component. This reduces NBR_{SLM} and, consequently, lowers the overall quality of the estimate.

In summary, the choice of the BP extraction method should be guided by the available time, resources, and the number of images. The SLM-based extraction method is ideal when time or resources are limited, as it requires fewer images (no more than 10) and a simpler setup. For optimal results in this mode, an ISO setting between 100 and 200 is recommended. Conversely, for the highest BP quality, the NL-based extraction method is preferred, although it demands more controlled conditions and at least 50 images in portrait mode. To achieve the best results under the NL case, we recommend maintaining a constant luminance level between 150 and 200, ensuring no color channel saturation, and to use as large a value of L as possible. While the preceding discussion has focused on a particular base pattern $\mathbf{P}$ representing the SDNP (specifically the one embedded by the iPhone 15 with a 24MP resolution), similar conclusions can be drawn from other resolutions of the pattern (i.e., 12MP and 7MP) and also distinct base patterns identified and discussed in [1, Sect. V].

4.1 Apple’s BP Across Resolutions and Aspect Ratios

The new 24MP resolution starting with the iPhone 15 embeds a BP that closely resembles the one used at 12MP.

Change of Resolution	Aspect Ratio	Image Size ($H' \times W'$)	Shifts
24MP $\left(\begin{smallmatrix} H = 5712 \\ W = 4284 \end{smallmatrix}\right) \rightarrow$ 18MP	16:9	5712×3213	$i_0 = 0, j_0 = \lfloor \frac{W}{8} \rfloor + 1$
	1:1	4284×4284	$i_0 = \frac{H-W}{2}, j_0 = 0$
12MP $\left(\begin{smallmatrix} H = 4032 \\ W = 3024 \end{smallmatrix}\right) \rightarrow$ 9MP	16:9	4032×2268	$i_0 = 0, j_0 = \frac{W}{8}$
	1:1	3024×3024	$i_0 = \frac{H-W}{2}, j_0 = 0$
7MP $\left(\begin{smallmatrix} H = 3088 \\ W = 2316 \end{smallmatrix}\right) \rightarrow$ 5MP	16:9	3088×1738	$i_0 = 0, j_0 = \lfloor \frac{W}{8} \rfloor - 1$
	1:1	2316×2316	$i_0 = \frac{H-W}{2}, j_0 = 0$

Table 4: Support for Distinct Aspect Ratios

By measuring the NCC between $\hat{\mathbf{P}}_{\text{NL}}$ at 12MP and a resized version of the 24MP $\hat{\mathbf{P}}_{\text{NL}}$ (downsampled using a bilinear kernel) we obtain a value of 0.91, indicating that the underlying structure is very similar. This suggests that the BP extracted from 24MP images could be used to detect its presence in 12MP portrait images (and vice versa), provided the patterns are correctly scaled spatially. However, for optimal matching it is generally preferable to use a reference BP extracted directly from images at the target resolution (either 12MP or 24MP). Notably, this similarity does not apply to the BP embedded at 7MP. Even after an exhaustive search across different crops, shifts and resizing factors, we were unable to match the 7MP BP to the 24/12MP BP. This indicates that a different BP (or possibly a slightly altered version generated with a different seed) may be used for the front camera, which captures images at 7MP resolution.

When selecting aspect ratios other than the default 4:3, such as 1:1 (square) or 16:9, a distinct BP is not embedded. Instead, a cropped version of the 4:3 reference BP is used. The samples from the $H \times W$ reference BP $\mathbf{P}$ in the 4:3 format (for resolutions 24MP, 12MP, or 7MP) are cropped to generate the $H' \times W'$ reference BP $\mathbf{P}^{(\text{AR})}$, where AR stands for Aspect Ratio (1:1 or 16:9), for the other aspect ratios, resulting in 18MP, 9MP, or 5MP resolution, respectively. The elements of the cropped pattern are given by

$$P_{i,j}^{(\text{AR})} = P_{i+i_0, j+j_0},$$

for all $i = 0, \dots, H' - 1, j = 0, \dots, W' - 1$, and i_0 and j_0 are the vertical and horizontal shifts, respectively, to align the pattern with the desired aspect ratio (AR), as specified in Table 4. Note that if the image is taken in landscape orientation, i.e., $W > H$, the values of i_0 and j_0 must be swapped, replacing also $\frac{W-H}{2}$ by $\frac{H-W}{2}$, and $\frac{H}{8}$ by $\frac{W}{8}$ in the rightmost column of Table 4.

4.2 Other Apple’s BP features

As part of the characterization of the extracted BPs, we also analyzed them in the frequency domain. Figure 9a shows the Power Spectral Density (PSD) of $\hat{\mathbf{P}}_{\text{SLM}}$ extracted from HEIF format images at 12MP resolution




Figure 9: PSD of the extracted BP $\hat{\mathbf{P}}_{\text{SLM}}$ from HEIF format images from the iPhone 12 mini (a). Autocorrelation of $\hat{\mathbf{P}}_{\text{SLM}}$ limited to $k, l \in [-5, 5]$ (b). Autocorrelation of $\hat{\mathbf{P}}_{\text{SLM}}$ removing central values for $k, l \in [-5, 5]$ and limiting visualization to $k, l \in [-110, 110]$ (c).

from the iPhone 12 mini (a similar result was obtained for the iPhone 15). As seen in the figure, most of the BP energy is concentrated in the low frequencies, and the nearly perfect circular shape likely indicates post-processing involving a low-pass filter applied to the noise pattern. Regarding the autocorrelation of $\hat{\mathbf{P}}_{\text{SLM}}$ for lags k and l (limited to $k, l \in [-5, 5]$), shown in Figure 9b, it is characteristic of a colored (as opposed to white) process and consistent with low-pass filtering. Moreover, after removing values around $(k, l) = (0, 0)$, two peaks located at $(k, l) = (-100, -100)$ and $(k, l) = (100, 100)$ become visible in Figure 9c. This specific feature is not present in BPs extracted from other iPhones we analyzed, such as the iPhones 13, 14, and 15, but has appeared in others, like the iPhone 12 Pro. While we suspect this feature may be linked to the specific BP generation process, its presence suggests a repetition of the pattern, which will be further analyzed in future work.

The extracted BPs share some similarities with a low-pass filtered white noise process (as hinted by Figs. 9a and 9b). In Figure 10a, we depict the upper-left 128×128 patch of $\hat{\mathbf{P}}_{\text{SLM}}$ extracted from HEIF format images at 12MP resolution on the iPhone 15. To analyze the grain size distribution in the extracted BP at a given resolution, we applied granulometry, which involves performing morphological openings with structuring elements of varying radii. This process enhances BP contrast before computing the sum of pixel values in the opened images for dif-




Figure 10: Upper-left 128×128 patch of $\hat{\mathbf{P}}_{\text{SLM}}$ extracted from HEIF format images at 12MP resolution on the iPhone 15 (a). Granulometry analysis of $\hat{\mathbf{P}}_{\text{SLM}}$ (b). Upper-left 600×600 block from $\hat{\mathbf{P}}_{\text{NL}}$ (iPhone 12 mini) that highlights the traces from Apple’s tiling artifact zooming at the intersection of the 512×512 grid.

ferent radii. The derivative of these sums is then plotted in Figure 10b, illustrating BP granulometry at each resolution. For comparison, we also include the same analysis for a white Gaussian noise process of the same size as the 12MP resolution. As shown, the BP exhibits a grain size extending up to 2 pixels, consistent with the autocorrelation support that is observed in Figure 9b. In contrast, the white Gaussian noise has a grain size limited to 1 pixel.

Finally, focusing on the extraction of $\hat{\mathbf{P}}_{\text{NL}}$ from HEIF format images at 12MP resolution on the iPhone 12 mini (with $K = 3$ for the denoising filter), we observed a grid pattern with a periodicity of 512×512 samples. Figure 10c shows the upper-left 600×600 block of the extracted BP where a zoom is made at the grid intersections. Initially, we believed this artifact was part of the BP. However, after reviewing the work in [4], we realized that this pattern is actually the tiling applied by Apple in the HEIF format. Interestingly, this tiling artifact is not visible in images captured by the iPhone 15, suggesting it may be specific to certain iPhone models.

4.3 Variability of $g(\cdot)$ across different devices

In this section, we investigate the variability of the estimated brightness-dependent function $g(\cdot)$, aiming to determine whether it remains consistent across different iPhone models, devices, and iOS versions, or instead exhibits deviations that could provide forensic clues about the specific capture conditions. In [1, Sect. IV-B], only portrait images from our iPhone 15 (D81 in Tab. 6) were analyzed. However, we observed that the curves obtained for ISO 125 and ISO 200 (captured under iOS 17 and iOS 18, respectively) do not show notable differences, as illustrated in Fig. 11a.

Following the same procedure described in [1, Sect. IV-B], we acquired a new set of NL portrait images using a different iPhone 15 (same model, distinct device, referred to as iPhone 15[†]). The captures were taken with a ColorChecker-like board displayed on a screen behind the subject, generating uniform gray patches with the SDNP

present in the background (see left panel of Fig. 12a). A total of $T = 86$ images were captured at ISO 250 and 24 MP resolution. From each image, $M = 12$ co-located $B \times B$ patches ($B = 512$ pixels) were extracted (highlighted in red in the right panel of Fig. 12a), yielding $M \cdot T = 1032$ distinct patches. The resulting empirical pairs $(z_{m,t}, \hat{g}(y'_{m,t}))$ for the luminance channel are plotted in Fig. 11b, along with the curve obtained from our iPhone 15 (D81) at ISO 125 for comparison. Despite the absence of samples near extreme luminance values (close to 0 and 255) for iPhone 15[†], both curves exhibit strong consistency (apart from minor deviations in the $[5, 15]$ range) suggesting that the brightness-dependent function remains stable across devices of the same model.

To evaluate cross-model consistency, we repeated the experiment with an iPhone 16 (D88), capturing $T = 97$ images at ISO 200 and processing $M = 14$ co-located blocks (see Fig. 12b). The resulting curve is compared with that from our iPhone 15 (D81) at the same ISO in Fig. 11c. Although the two estimated functions are not identical, they show a strong resemblance. We hypothesize that the small discrepancies may result from bias introduced by variability in the noise variance within the model described in Sect. IV-B, which can affect the estimation of $\hat{g}(\cdot)$.

Finally, the parameter inducing the largest variability in the estimated brightness-dependent function $g(\cdot)$ was the ISO setting, which is also partially influenced by lighting conditions during capture. The strongest variations were observed in portrait images from the iPhone 16 (D88), particularly between ISO 200 and ISO 500, as shown in Fig. 11d. We plan to further investigate these differences, which may stem from contrast adjustments applied after the bokeh effect or from variability in the noise variance. As future work, we aim to explore how the knowledge of $\hat{g}(\cdot)$ could be leveraged to enhance BP extraction and detection.

5 Expanded Experimental Results

5.1 Threshold Selection for BP Detection

As described in [1, Sect. VI-A], Apple portrait images are detected when the NCC between the noise residue $\mathbf{W}$ of the analyzed image and a reference BP estimate $\hat{\mathbf{P}}$ exceeds a predefined threshold β , i.e., $\rho(\mathbf{W}, \hat{\mathbf{P}}) > \beta$. Our threshold selection strategy aims to minimize cross-device misattributions caused by Apple’s SDNP embedded in portrait mode images and is therefore designed to achieve a very low False Positive Rate (FPR), e.g., $\text{FPR} = 10^{-3}$. To this end, we collected a large set of negative samples (images captured without Apple portrait mode) at 12MP resolution. From the dataset used in [2], we selected a total of 9601 images from various Apple and Samsung devices sharing the same 12MP resolution, as summarized in Tab. 5. Following the procedure in [1, Sect. VI-A], 7303




Figure 11: Estimated brightness-dependent function $\hat{g}(\cdot)$ across different iPhone models, devices, iOS versions, and ISO settings. Comparison between iOS 17 and iOS 18 using portrait images from iPhone 15 (D81) (a), showing consistent behavior across versions. Comparison between two iPhone 15 devices (D81 and iPhone 15[†]) captured respectively at ISO 125 and ISO 250 (b), demonstrating strong inter-device consistency. Comparison between iPhone 15 (D81) and iPhone 16 (D88) at ISO 200 (c), showing minor deviations likely due to noise variance. Variability observed in iPhone 16 (D88) across ISO 200, 320, and 500 (d), where the strongest differences in $\hat{g}(\cdot)$ are noted.




Figure 12: Examples of portrait images of the ColorChecker scene captured at full resolution (24MP). The left panel (a) displays two samples from a set of $T = 86$ images taken at ISO 250 with the iPhone 15[†], while the right panel (b) shows two samples from a set of $T = 97$ images taken at ISO 200 with the iPhone 16 (D88). The processed $B \times B$ blocks ($B = 512$) are highlighted in red in the rightmost image of each case.

images were randomly selected and processed (reserving 2298 to match the number of positive samples, thus obtaining a balanced test set as in [1, Sect. VII-A]). For each image, the NCC was computed across multiple BP versions (eight in total, as detailed in [1, Sect. V]) and all possible 90° rotation variants, taking the maximum NCC value as the final score compared against β .

The histogram of the resulting maximum NCC values $\rho_{\max}$ is shown in Fig. 13a, where three thresholds are marked: 0.0031 for $\text{FPR} = 10^{-2}$, 0.0072 for $\text{FPR} = 10^{-3}$, and 0.011 for $\text{FPR} = 10^{-4}$. Note that the threshold value strongly depends on the BP estimates used. In previous experiments, the threshold corresponding to $\text{FPR} = 10^{-2}$ was approximately 0.01, but after incorporating new and improved BP estimates (see [1, Sect. V] for details), this value decreased to 0.0031. Based on these results, we selected the threshold corresponding to $\text{FPR} = 10^{-3}$, i.e., $\beta = 0.0072$, as the predefined threshold for BP detection.

The detector is then evaluated on the balanced testing set (as also reflected in [1, Sect. VII-A]), consisting of the remaining 2,298 negative samples from Tab. 5




Figure 13: Histogram of the maximum NCC values $\rho_{\max}$ obtained after processing the set of negative samples (i.e., images captured without Apple portrait mode), highlighting three different thresholds are depicted for distinct FPRs: 10^{-2} , 10^{-3} , and 10^{-4} (a). ROC curve obtained on the test dataset including negative (2,298) and positives (2,298) samples from Tabs. 5 and 6, respectively; achieving an AUC of 0.998 (b).

and an equal number of portrait images (2,298) from Tab. 6. The resulting ROC curve is presented in Fig. 13b, with a zoomed-in view highlighting the detector’s excellent performance, achieving an AUC of 0.998. Using the predefined threshold $\beta = 0.0072$, the detector achieves $\text{FPR} = 4.3 \cdot 10^{-4}$, in line with the target $\text{FPR} = 10^{-3}$, and $\text{TPR} = 0.9761$, representing very high detection accuracy at a minimal false alarm rate. For comparison, using a threshold of 0.0031 yields $\text{FPR} = 9.1 \cdot 10^{-3}$ and $\text{TPR} = 0.9909$, indicating near-perfect detection performance.

5.2 Threshold Selection for BP Localization

The strategy used to determine the threshold α for deriving the binary mask $\mathbf{M}^{(\text{PRNU})}$ (as described in [1, Sect. VI-B]) involves capturing a set of 77 images with a subject positioned in the lower part of the frame (similar

Table 5: List of smartphones employed to build the negative sample set, indicating brand, model, number of devices, and total images. All images have a 12MP resolution and were sourced from the dataset used in [2].

Brand	Model	#devices	#images
Apple	iPhone 7	10	693
Apple	iPhone 7 Plus	7	431
Apple	iPhone 8	9	622
Apple	iPhone 8 Plus	10	628
Apple	iPhone X	6	322
Apple	iPhone XR	5	348
Apple	iPhone XS	7	438
Apple	iPhone XS Max	9	589
Apple	iPhone 11	7	457
Apple	iPhone 11 Pro	7	304
Apple	iPhone 11 Pro Max	9	495
Samsung	SM-A505U	3	209
Samsung	SM-G930V	8	545
Samsung	SM-G935F	7	490
Samsung	SM-G950U	8	511
Samsung	SM-G955U	10	700
Samsung	SM-G960U	7	490
Samsung	SM-G965U	9	516
Samsung	SM-G970U	4	280
Samsung	SM-G973U	6	304
Samsung	SM-G975U	8	461

to the example in [1, Fig. 3]) so that the upper half of each image contains Apple’s bokeh effect and thus the SDNP. Unlike the approach used for obtaining the NL-based BP estimate, these images feature non-uniform backgrounds, as our goal is to localize the BP under diverse conditions. From this set, positive samples are created by cropping distinct $B \times B$ blocks from the upper half of the images ($0 \leq i \leq H/2 - 1$). Using $B = 21$ (the value employed for BP localization), we obtain over one million blocks for the positive class. For simplicity, we use the corresponding unmodified originals (i.e., images without the bokeh effect) to generate the negative class blocks.

For each block, the NCC between its residue and the BP estimate (specifically $\hat{\mathbf{P}}_{\text{NL}}(\text{Hi})$ obtained from our iPhone 15, as discussed in Sect. 4) is computed. The resulting histogram for a subset of 10^6 blocks per class is shown in Fig. 14a, revealing a noticeable overlap between the two distributions. This overlap arises partly from the small block size (reducing discriminative power) and partly from Apple’s refinements to SDNP embedding, which make it more selective by excluding certain regions (e.g., edges, as discussed in [1, Sect. VII-C]). Consequently, some 21×21 patches cover edge areas where the SDNP signal is weak or absent. When using larger blocks ($B = 256$), the overlap becomes negligible, as illustrated in Fig. 14b.

Since our objective is to achieve fine-grained segmenta-




Figure 14: Histograms of ρ values obtained for positive and negative samples used to estimate the threshold α for BP localization. Results for $B = 21$ (a), showing a noticeable overlap between classes due to the reduced block size and the selective nature of Apple’s SDNP embedding. Results for $B = 256$ (b), where the overlap becomes negligible, confirming the stronger separability of larger blocks.

tion of SDNP-affected regions (so that $\mathbf{M}^{(\text{PRNU})}$ primarily includes areas with strong underlying PRNU), we set the threshold α at the Equal Error Rate (EER) operating point (i.e., $\text{EER} = 0.1322$). This yields $\alpha = 0.07$, which, when applied to the remaining nearly one million samples per class, results in $\text{FPR} = 0.1302$ and $\text{TPR} = 0.8625$. Although this is not the optimal operating point, as discussed in [1, Sect. VII-B], we believe that developing an adaptive strategy to estimate α on a per-image basis could substantially improve BP localization performance. However, this remains an open problem for future research.

5.3 On Replicating Baracchi *et al.*’s Results

In validating our BP-aware PRNU matching approach (described in [1, Sect. VI-B]), we compared it against the state-of-the-art method proposed by Baracchi *et al.* [8]. However, as noted in [1, Sect. VII-C], we encountered difficulties reproducing their results. Upon further investigation, we suspect that changes in Apple’s portrait rendering pipeline over time have reduced the effectiveness of using depth map information alone to exclude regions affected by the SDNP. This may explain the discrepancies observed. We reached out to the authors of [8], who kindly shared their experimental data. Although the original code, PRNU extraction data, and exact image sets for positive and negative classes were not available, we made our best effort to replicate their results.

In [8, Sect. 4.2], the authors use 50 images to extract the PRNU fingerprint of an iPhone X (iOS 12.1.4) and 100 portrait images from the same device for the positive class. However, the shared dataset we received included only 140 images in total. Therefore, in our experiments, we used 40 images for PRNU extraction and the remaining 100 portrait images as the positive class.

Following the description in [8, Sect. 4.2], we constructed the negative class using 100 portrait images from

Table 6: List of smartphones used to capture the 12MP portrait images, indicating ID, model, iOS version, number of images, and data source. Portrait images used exclusively for BP extraction and not for BP detection are also listed.

ID	Model	iOS	#images	Source	ID	Model	iOS	#images	Source
D01	iPhone 7 Plus	10.3.2	19	[2]	D48	iPhone 12 mini	17.4.1,	27	[5]
D02	iPhone 7 Plus	10.3.1,	311	[5]			17.5.1		
		10.3.2,			D49	iPhone 12 mini	16.5.1	11	[5]
		10.3.3,11.2.6,			D50	iPhone 12 mini	14.1	8	[6]
		11.3, 11.4			D51	iPhone 12 mini	15.1,	110	Ours
D03	iPhone 7 Plus	11.0.3	46	[5]			15.1.1,		
D04	iPhone 7 Plus	10.3.1,	12	[5]			15.2,		
		10.3.3,					15.2.1,		
		11, 11.2.2					16,		
D05	iPhone 7 Plus	11.4	13	[5]			16.5.1,		
D06	iPhone 7 Plus	10.2.1,	118	[5]			16.6,		
		10.3,					16.6.1,		
		10.3.2					17.4.1		
D07	iPhone 7 Plus	10.1	2	[6]	D52	iPhone 12 Pro	14.1	18	[6]
D08	iPhone 8 Plus	11.0.2	11	[2]	D53	iPhone 12 Pro	17.1.1	20	[9]
D09	iPhone 8 Plus	11.3.1	16	[2]	D54	iPhone 12 Pro Max	14.1	17	[6]
D10	iPhone 8 Plus	11.4,	61	[5]	D55	iPhone 13	16.1	100	[5]
		11.4.1			D56	iPhone 13	15	4	[6]
D11	iPhone 8 Plus	11.2.5	17	[5]	D57	iPhone 13	17.1.1	20	[9]
D12	iPhone 8 Plus	11.0.1	13	[6]	D58	iPhone 13	17.1.2	20	[9]
D13	iPhone X	13.4.1	8	(C19) [7]	D59	iPhone 13	17.1.1	20	[9]
D14	iPhone X	11.1.2,	21	[8]	D60	iPhone 13	16.1.1	18	[9]
		11.2			D61	iPhone 13	17.1.1	20	[9]
D15	iPhone X	12.2	43	[8]	D62	iPhone 13 mini	15.0.1	2	[6]
D16	iPhone X	12.1.4	140	[8]	D63	iPhone 13 mini	17.2	20	[9]
D17	iPhone X	11.4,	22	[5]	D64	iPhone 13 Pro	16.6	76	[5]
		11.4.1			D65	iPhone 13 Pro	16.3.1	10	[5]
D18	iPhone X	11.2.6	44	[5]	D66	iPhone 13 Pro	15.0.1	12	[6]
D19	iPhone X	11.3.1	23	[5]	D67	iPhone 13 Pro Max	15	14	[6]
D20	iPhone X	11.0.1,	9	[6]	D68	iPhone 13 Pro Max	17	20	[9]
		11.1			D69	iPhone 13 Pro Max	17.1.1	13	[9]
D21	iPhone X	16.6	24	[9]	D70	iPhone SE (3rd)	15.4	6	[6]
D22	iPhone XR	12.2,	9	[2]	D71	iPhone 14	16.0.1	6	[6]
		12.3.1			D72	iPhone 14	17.1.2	20	[9]
D23	iPhone XR	12.1.2	15	[8]	D73	iPhone 14	17.1.2	20	[9]
D24	iPhone XR	12.0.1	5	[6]	D74	iPhone 14 Plus	16.0.2	8	[6]
D25	iPhone XS	12	11	[6]	D75	iPhone 14 Pro	16.0.1	14	[6]
D26	iPhone XS Max	12.2	10	[8]	D76	iPhone 14 Pro Max	17.0.3	13	[9]
D27	iPhone XS Max	12	12	[6]	D77	iPhone 14 Pro Max	17.4.1	4	[5]
D28	iPhone 11	13.2	11	[2]	D78	iPhone 14 Pro Max	16.0.1	18	[6]
D29	iPhone 11	13.1.2	3	[2]	D79	iPhone 14 Pro Max	17.1.2	20	[9]
D30	iPhone 11	13.4.1	10	(C20) [7]	D80	iPhone 14 Pro Max	17.1.2	14	[9]
D31	iPhone 11	13.3.1	10	(C21) [7]	D81	iPhone 15	17, 18, 26	113	Ours
D32	iPhone 11	16.6.1	5	[9]	D82	iPhone 15	17.4.1,	5	[5]
D33	iPhone 11	13.1	9	[6]			17.5.1		
D34	iPhone 11	17.1.1	20	[9]	D83	iPhone 15	17, 17.0.2	17	[6]
D35	iPhone 11 Pro	13.2.3,	29	[2]	D84	iPhone 15	17.0.3	2	[9]
		13.3.1			D85	iPhone 15 Plus	17.0.2	10	[6]
D36	iPhone 11 Pro	13.2	62	[2]	D86	iPhone 15 Pro	17.0.2	18	[6]
D37	iPhone 11 Pro	13.1.3	28	[2]	D87	iPhone 15 Pro Max	17.0.2	18	[6]
D38	iPhone 11 Pro	17.6.1	4	[9]	D88	iPhone 16	18	50	Ours
D39	iPhone 11 Pro Max	13.4	12	(C22) [7]	D89	iPhone 16	18	8	[6]
D40	iPhone 11 Pro Max	13.1	13	[6]	D90	iPhone 16 Plus	18	8	[6]
D41	iPhone 11 Pro Max	16.7.2	20	[9]	D91	iPhone 16 Pro	18	9	[6]
D42	iPhone SE (2nd)	13.4.1	4	[6]	D92	iPhone 16 Pro Max	18	6	[6]
D43	iPhone SE (2nd)	16.1.2	20	[9]	D93	iPhone 16e	18.3.1	8	[6]
D44	iPhone 12	17.1.2	52	[5]	D94	iPhone 17	26	8	[6]
D45	iPhone 12	14.4.2	6	[5]	D95	iPhone 17 Pro	26	12	[6]
D46	iPhone 12	14.1	7	[6]	D96	iPhone 17 Pro Max	26	8	[6]
D47	iPhone 12 mini	16.6	3	[5]	D97	iPhone Air	26	8	[6]




Figure 15: Panel (a) shows the binary mask used by Baracchi *et al.* (left) and the corresponding per-pixel contribution map excluding likely SDNP regions (right) from the first experiment. Panel (b) presents the same results for the second experiment, using a different reference set for PRNU extraction.

PRNU-based method	Baseline	Weighted	Binary	BP-aware
1 st experiment	0.5925	0.4357	0.0246	0.0064
2 nd experiment	0.6466	0.5226	0.0423	0.0074

Table 7: NCC values between the BP used in the corresponding portrait images and the extracted PRNUs using each of the four PRNU-based methods: baseline, weighted [8], binary [8], and our BP-aware approach [1]. The minimum NCC value achieved for each experiment is highlighted in bold.

the shared dataset, which includes images from devices similar to the one used in the positive class. Specifically, the negative set consists of 10 images from an iPhone XS Max, 15 from an iPhone XR, and 75 from three different iPhone X devices (43, 21, and 11 images, respectively). Notably, the 21 images from one of these iPhone X devices were captured with older iOS versions (11.1.2 and 11.2). As a result, these images embed a different BP (matching that of the iPhone 8 Plus with iOS 11.3.1, as shown in [1, Tab. 1]) and differ from the BP of the positive class, which aligns with that of the iPhone XR running iOS 12.3.1 (also in [1, Tab. 1]). As such, these samples are unlikely to significantly contribute to PRNU collisions. In this setup, our BP-aware approach is applied with the two identified BPs, using parameters $K = 5$ and $\alpha = 0.07$.

We conducted two experiments, each using a different partition of the available 140 images, allocating 40 for PRNU extraction and 100 for the positive testing class, and evaluated four methods: the baseline, the weighted and binary approaches from [8], and our BP-aware method [1]. The goal was to assess whether the size of the mask obtained by the binary method from Baracchi *et al.* influences the final results. The masks generated by the binary method in each experiment are shown in Fig. 15, along with the sum of all masks $\mathbf{M}^{(\text{PRNU})}$ (defined in [1, Sect. VI-B]) for our BP-aware approach as a reference. As observed, the binary mask from the second experiment is smaller compared to the first.

In a similar way, Table 7 presents the NCC values computed between the BP in the corresponding portrait images and the extracted PRNUs using each of the four PRNU-based methods. The results show that BP leakage




Figure 16: ROC curves comparing PRNU-based source camera verification performance across different methods: baseline, weighted [8], binary [8], and our BP-aware approach [1]. Each experiment uses a different reference set of 40 portrait images for PRNU extraction.

into the PRNU is generally higher in the second experiment, despite the smaller binary mask. Additionally, the NCC values progressively decrease from the baseline to the weighted and binary methods by Baracchi *et al.*, with our BP-aware approach achieving the lowest BP leakage.

The ROC curves obtained from both experiments are shown in Fig. 16. We were able to closely replicate the performance of the weighted and binary methods proposed by Baracchi *et al.* (cf. [8, Fig. 8]); however, the baseline results differ markedly from those reported in [8], and we are unable to explain this discrepancy. Comparing the two experiments, the baseline method shows little variation, while the weighted and binary approaches of Baracchi *et al.* perform noticeably better in the first experiment (likely due to reduced BP leakage, as observed in Table 7). In contrast, our BP-aware method delivers stable performance across both experiments and consistently achieves the best overall results.

5.4 Forgery Localization Using BP-aware PRNU Analysis

After analyzing the BP-driven NCC maps in [1, Fig. 10], one might consider using the BP as a known reference in Apple portrait images. Any manipulation affecting a BP-embedded region would become evident through the BP-driven NCC map. For instance, Figure 17c shows such a map for the manipulated image in Figure 17b, which was altered using Google Photos AI editor to remove content via AI inpainting. This image is a modified version of a portrait from [7] (Figure 17a). As expected, the edited region exhibits a clear lack of correlation with the BP.

However, the absence of correlation alone does not indicate whether Apple’s algorithm intentionally omitted the BP or if the image has been tampered with. To address this, we incorporate the PRNU of the capturing device whenever available. In this case, we estimate the PRNU




Figure 17: Forgery localization using BP- and PRNU-driven NCC maps. Original portrait images from [7] (a, f). Manipulated version edited with Google Photos AI inpainting (b, g). Overlaid BP-driven NCC map (c, h). PRNU-driven NCC map (d, i). Combined forgery map (e, j).

from the “flat” and “nat” images of the C21 device in the [7] dataset. We then compute a PRNU-driven NCC map, similar to the BP-driven version described in [1, Sect. VI-B], but using a different block size (32×32) and a larger neighborhood (512×512) to enhance PRNU matching, given its lower energy compared to the BP. Figure 17d shows the PRNU-driven NCC map for the manipulated image (Figure 17b). Although it lacks the fine-grained detail of the BP-driven map, correlation loss is (to some extent) still evident in the tampered region.

To carry out forgery localization, we propose combining the two maps. First, we normalize each by dividing pixel values by the map’s maximum value. Then, we sum the normalized maps and visualize the result using a heatmap, where lower values appear as hotter colors. The resulting forgery map (Figure 17e) highlights the manipulated area. However, some artifacts unrelated to the forgery appear, primarily due to noise in the PRNU map. A more accurate PRNU estimate, obtained from a larger set of images, could enhance localization. Additionally, the manipulated image underwent a complex post-processing chain (including recompression upon upload to Google Photos, global editing via the AI tool, and final compression when saving) likely diminishing PRNU traces and introducing noise.

Despite these challenges, this proof of concept shows that forgery localization is feasible in real-world scenarios when the manipulated region originally contained Apple’s BP and a PRNU estimate is available. However, further analysis is needed to minimize false positives in genuine Apple portrait images, which we plan to explore in future work. Another example of an edited Apple portrait image

appears in the bottom panel of Figure 17.

5.5 Detecting and Identifying Apple’s BP “in the wild”

After successfully detecting and identifying Apple’s BP under controlled conditions (see [1, Sect. VII-A]) and following the methodology applied in [10] with the Adobe pattern, we extended our study to assess the performance of our approach in an uncontrolled open-set scenario using the FFHQ (in-the-wild-images) dataset [11]. This dataset consists of 70,000 images of human faces, crawled from Flickr. While the dataset provides URLs to the original Flickr images, all collected images are stored in PNG format. To better simulate a real-world forensic scenario (where information about the original source might not be available) we conducted our analysis directly on the PNG images. Given the large dataset size and the wide variety of image resolutions, we restricted our study to 12MP images (i.e., those with a resolution of 4032×3024 or its landscape version), resulting in a total of 427 images (actually, 449 images, but 22 are duplicates).¹

For BP detection and identification, we used the 8 available BP estimates (see [1, Sect. V]). In all cases, NL-based BPs $\hat{\mathbf{P}}_{\text{NL}}$ were employed, except for BPs ⑤ and ⑦, where SLM-based BPs $\hat{\mathbf{P}}_{\text{SLM}}$ were used. Although the dataset was published in early 2019, we included all newer BPs (i.e., ⑤–⑦) to assess whether patterns generated by post-2019 devices could influence detection or

¹To establish the ground truth for portrait images, we downloaded the original versions of the 427 files, successfully recovering 376. For the remaining 51, which were no longer available on Flickr, metadata could not be retrieved, preventing proper labeling.

identification performance.

Using the predefined threshold $\beta = 0.0072$ (see Sect. 5.1), we obtained a total of 96 positive detections out of 427 images. Considering only the 376 images with available ground-truth labels (73 positives and 303 negatives), 78 images were classified as positives, of which 72 correspond to true Apple portrait images, yielding a True Positive Rate (TPR) of 0.9863. Among the six false positives (corresponding to a 0.0198 False Positive Rate, FPR), 1 originated from an iPhone 7 image with heavy white saturation (indicating that it might have been post-processed), while the remaining five were DSLR images captured with an Olympus E-30 and post-processed using Apple’s *Aperture* software (as indicated by EXIF meta-data).²

Regarding misdetections, only one portrait image (captured with an iPhone 7 Plus) was missed, likely due to post-processing with the *VSCO* photo editor (as indicated by EXIF metadata), which may have removed BP traces. For BP identification, only one incorrect attribution was observed: an image captured with an iPhone 7 Plus (with EXIF tag Software: *Photos 2.0*) matched BP ⑧, suggesting that Apple’s *Photos* app may embed its own BP-like patterns (as hinted in [1, Sect. V]). Although the similarity score (0.0743) exceeded the threshold, the corresponding BP-driven NCC map (see Fig. 18) displays a characteristic arc pattern in the top-left corner, resembling the one reported in the top-right corner of [1, Fig. 10a], reinforcing the hypothesis that Apple BPs originate from a specific algorithm. Future work will further explore BPs arising from Apple software such as *Photos* and the aforementioned *Aperture*.

For the 51 images lacking ground-truth labels, 18 were detected as positives. All 18 exhibit BP-driven NCC maps consistent with Apple portrait images, supporting their classification as true positives. Figure 19 illustrates two such examples. The NCC map in Fig. 19a, obtained using BP ③, indicates that the image was likely captured with an iPhone 8 Plus or iPhone X running iOS 11 (see [1, Tab. 1]), while Fig. 19b, derived from BP ②, suggests capture by an iPhone 7 Plus running iOS 11. In total, 90 portrait images (72 labeled and 18 unlabeled) were positively identified, representing 21.08% of the 12MP images in the FFHQ (in-the-wild) dataset. This proportion indicates that the presence of Apple’s BP in real-world images is far from negligible.

5.6 Additional details on BP Detection under Post-Processing

Using the same originally captured images as in [1, Sect. VII-D], we further assess the robustness of our BP

²We extracted a pattern from the *Aperture* software correlating with BP ③ (NCC = 0.0158). Further investigation is needed to understand the origin of these shared features.




Figure 18: Example of misattributed BP. The image, captured with an iPhone 7 Plus (Software: *Photos 2.0*), was incorrectly matched to BP ⑧ (a pattern only observed on iPhone 8 Plus and iPhone X models). The BP-driven NCC map reveal characteristic arc patterns, similar to those shown in [1, Fig. 10a], suggesting that Apple’s *Photos* app may embed its own BP-like pattern, consistent with the hypothesis that Apple BPs stem from a specific algorithm.




Figure 19: BP-driven NCC maps for images with no ground truth labeling from FFHQ. Portrait image whose BP matches BP with index ③ (i.e., it probably comes from an iPhone 8 Plus or an iPhone X according to [1, Tab. 1]) (a). Portrait image whose BP matches BP with index ② (i.e., it probably comes from an iPhone 7 Plus running iOS 11 according to [1, Tab. 1]) (b).

detection by testing various combinations of scaling factors (SF: 1, 0.75, 0.5, 0.25, 0.1) and JPEG quality factors (QF: N/A (no compression), 95, 90, 85, 80, 75, 70, 60, 50). The AUC values, summarized in Table 8, indicate that detection remains perfect across all QF values as long as $SF \geq 0.5$. However, performance drops significantly when SF is reduced to 0.1, with further degradation as compression increases (lower QF). These findings confirm that our approach remains highly robust across most practical forensic scenarios.

References

- [1] D. Vázquez-Padín, F. Pérez-González, and P. Pérez-Miguelé, “Apple’s Synthetic Defocus Noise Pattern: Characterization and Forensic Applications,” *IEEE Transactions on Information Forensics and Security*, pp. 1–16, 2026.

SF \ QF	N/A	95	90	85	80	75	70	60	50
1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
0.75	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
0.50	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
0.25	1.00	1.00	1.00	0.98	0.96	0.91	0.86	0.82	0.76
0.10	0.86	0.53	0.46	0.46	0.45	0.43	0.45	0.49	0.46

Table 8: AUC Values for Various Combinations of JPEG Quality Factors and Scaling Factors.

- [2] M. Iuliani, M. Fontani, and A. Piva, “A Leak in PRNU Based Source Identification. Questioning Fingerprint Uniqueness,” *IEEE Access*, vol. 9, pp. 52 455–52 463, 2021.
- [3] “Apple API: builtInTripleCamera.” [Online]. Available: <https://developer.apple.com/documentation/avfoundation/avcapturedevice/devicetype/3377621-builtintriplecamera/>
- [4] S. McKeown and G. Russell, “Forensic Considerations for the High Efficiency Image File Format (HEIF),” in *2020 International Conference on Cyber Security and Protection of Digital Services*, 2020, pp. 1–8.
- [5] [Online]. Available: <https://www.flickr.com/services/api/>
- [6] [Online]. Available: <https://www.gsmarena.com/>
- [7] C. Albisani, M. Iuliani, and A. Piva, “Checking PRNU Usability on Modern Devices,” in *2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 2535–2539.
- [8] D. Baracchi, M. Iuliani, A. G. Nencini, and A. Piva, “Facing Image Source Attribution on iPhone X,” in *Digital Forensics and Watermarking*, X. Zhao, Y.-Q. Shi, A. Piva, and H. J. Kim, Eds. Cham: Springer International Publishing, 2021, pp. 196–207.
- [9] [Online]. Available: <https://forms.gle/PJW8UbASFX5oBsUu9>
- [10] J. Butora and P. Bas, “Detection of the Adobe Pattern,” in *32nd European Signal Processing Conference*, Lyon, France, Aug. 2024.
- [11] T. Karras, S. Laine, and T. Aila, “A Style-Based Generator Architecture for Generative Adversarial Networks,” 2019. [Online]. Available: <https://arxiv.org/abs/1812.04948>