

A MODIFIED FOURIER-MELLIN APPROACH FOR SOURCE DEVICE IDENTIFICATION ON STABILIZED VIDEOS

Sara Mandelli*, Fabrizio Argenti[†], Paolo Bestagini*, Massimo Iuliani^{†‡}, Alessandro Piva^{†‡}, Stefano Tubaro*

*Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano - Milan, Italy

[†]Department of Information Engineering, University of Florence - Florence, Italy

[‡]FORLAB - Multimedia Forensics Laboratory - Prato, Italy

ABSTRACT

To decide whether a digital video has been captured by a given device, multimedia forensic tools usually exploit characteristic noise traces left by the camera sensor on the acquired frames. This analysis requires that the noise pattern characterizing the camera and the noise pattern extracted from video frames under analysis are geometrically aligned. However, in many practical scenarios this does not occur, thus a re-alignment or synchronization has to be performed. Current solutions often require time consuming search of the re-alignment transformation parameters. In this paper, we propose to overcome this limitation by searching scaling and rotation parameters in the frequency domain. The proposed algorithm tested on real videos from a well-known state-of-the-art dataset shows promising results.

Index Terms— Video forensics, sensor noise, Fourier-Mellin, PRNU, video stabilization

1. INTRODUCTION

Multimedia forensics keeps developing technologies to identify the camera originating a digital image or a digital video. Currently, the most promising technique to achieve this task is based on the analysis of Sensor Pattern Noise (SPN) noise or Photo Response Non-Uniformity (PRNU), left by the acquisition device into the visual content. This trace is useful to identify the video source since it is universal (i.e., every camera sensor introduces one) and unique (i.e., PRNU from two different sensors are uncorrelated) [1, 2]. Moreover, PRNU is proved to be significantly robust to commonly used processing, like JPEG compression [1], or uploading to social media platforms [3, 4].

PRNU based source identification process consists in verifying the match between a query image or video frame and a fingerprint characterizing a reference camera. The strategy involves two main steps: i) a reference fingerprint is derived from still images or videos acquired by the source device; ii) the query fingerprint is estimated from the investigated content and then compared with the reference to verify the possible match, in form of a correlation. If the query content was acquired by the reference camera, then a high correlation is expected.

This material is based on research sponsored by DARPA and Air Force Research Laboratory (AFRL) under agreement number FA8750-16-2-0173. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of DARPA and Air Force Research Laboratory (AFRL) or the U.S. Government.

The previous scheme works under the hypothesis of perfect geometrical alignment between the reference and test fingerprints. If a geometrical transformation is applied to the query content, a pixel grid misalignment between the query and the reference fingerprint arises, thus hindering the detection. Such a case occurs in multiple scenarios: when an image has been acquired with different resolution settings or it is cropped and resized due to the upload in a social media; if a malicious user slightly distorts an image to remove the sensor traces; when a query video is tested against a reference estimated from still images; when a video has been created in presence of electronic image stabilization. In all these cases, the PRNU extracted from the query is misaligned with the reference fingerprint, and thus a geometric re-synchronization between them has to be carried out before the matching operation.

The first solution to this problem was proposed in [5], where the case of cropped and downsampled images was studied. The authors show that it is possible to parameterize the Normalized Cross-Correlation (NCC) between the reference fingerprint and the query noise with respect to the scaling factor. The NCC peak position for a given scaling factor provides an estimate of the shift. While the NCC can be efficiently computed in the Fourier domain, a brute force search is needed to determine the scaling factor. By following the same rationale, more recent papers [6, 7, 8] extend the proposed methodology also considering rotation and the more challenging scenario of video analysis. As a matter of fact, modern acquisition pipelines usually include electronic stabilization that undermines PRNU-based attribution technique. In these cases, PRNU-based techniques only work if geometric transformations are properly estimated and compensated for, which is a computational complex operation.

In this paper, we focus on the problem of stabilized video sequences camera attribution based on PRNU. Specifically, we propose a method to align frame fingerprints with reference PRNU by recovering the scaling, shift and rotation parameters introduced by the electronic stabilization. We overcome the problem of computational complexity by searching for scaling and rotation parameters in the frequency domain thanks to a modified version of the Fourier-Mellin transform (FM). Results obtained on the well known Vision dataset [9] show that the proposed method provides extremely efficient results whenever rotation and scaling operations are applied to video frames. When also shift is taken into account, the gain compared against the state-of-the-art [8] depends on the video content.

2. BACKGROUND AND PROBLEM STATEMENT

In this section we introduce the background on Fourier-Mellin (FM) transform and define the problem we are tackling in this paper.

Fourier-Mellin Transform. The FM transform enables to estimate scale, rotation and shift transformations between two images in closed form [10].

Given an image \mathbf{I} , the FM transform is expressed as the log-polar mapping of the magnitude of the image Fourier transform, i.e.,

$$\text{FM}\{\mathbf{I}\} = \text{LP}\{|\mathbf{F}|\}, \quad (1)$$

where $\text{LP}\{\cdot\}$ is the operator computing log-polar mapping, and $|\mathbf{F}|$ is the magnitude of the Fourier transform.

Let us consider two images \mathbf{I}_a and \mathbf{I}_b that are linked through a similarity transformation, i.e., $\mathbf{I}_a = \mathcal{T}_{ab}\{\mathbf{I}_b\}$, where \mathcal{T}_{ab} applies the transformation identified by the matrix

$$\mathbf{T}_{ab} = \begin{bmatrix} s_{ab} \cdot \cos \alpha_{ab} & -s_{ab} \cdot \sin \alpha_{ab} & c_{abx} \\ s_{ab} \cdot \sin \alpha_{ab} & s_{ab} \cdot \cos \alpha_{ab} & c_{aby} \end{bmatrix}, \quad (2)$$

where s_{ab} represents scaling, α_{ab} rotation and $\mathbf{c} = [c_{abx}, c_{aby}]$ horizontal and vertical shift. In this scenario, it is possible to show that $\text{FM}\{\mathbf{I}_a\}$ is a shifted version of $\text{FM}\{\mathbf{I}_b\}$. More formally,

$$\text{FM}\{\mathbf{I}_a\}(\rho, \alpha) = \text{FM}\{\mathbf{I}_b\}(\rho - \log s_{ab}, \alpha - \alpha_{ab}), \quad (3)$$

where ρ is the radial coordinate and α the rotational coordinate. It is therefore possible to estimate scale s_{ab} and rotation α_{ab} by looking at the peak position of the phase correlation function between $\text{FM}\{\mathbf{I}_a\}$ and $\text{FM}\{\mathbf{I}_b\}$ [10] independently from shifts. Once s_{ab} and α_{ab} are estimated, the two images can be realigned apart from translation. The relative shift can then be estimated by looking at the peak position of the phase correlation computed between the two realigned images in the pixel domain [10].

Problem formulation. PRNU is typically modeled as a multiplicative noise pattern introduced by any device in all acquired images or videos [1, 11]. In the field of forensics analysis, it is well known that PRNU can be exploited for inferring whether an image was shot by a certain device. For instance, given a test image \mathbf{I} and a device PRNU \mathbf{K} , we can compute the Peak-to-Correlation Energy (PCE) between the noise residual \mathbf{W} extracted from the image and the PRNU pixel-wise scaled by \mathbf{I} , i.e., $\text{PCE}(\mathbf{W}, \mathbf{K} \cdot \mathbf{I})$. Indeed, PCE measures the correlation between the noise traces left on \mathbf{I} and the device PRNU independently of potential shift misalignment, as the correlation peak is searched over all possible mutual shifts between them. If PCE is greater than a confidence threshold, we attribute \mathbf{I} to the device [1, 11].

The extension of PRNU-based strategies for attributing video frames to a specific device suffers from some issues due, for instance, to higher compression rates and lower pixel resolutions. As a matter of fact, the previously described PCE test cannot be directly performed, being the PRNU resolution typically higher than the size of the recorded video frames. Moreover, in-camera video stabilization techniques, which are now becoming one of the must-have device specifications, strongly hinder the traces left by PRNU, as video frames may be warped by means of geometrical transformations (e.g., cropping, rotation, scaling, etc.) in order to generate a stable video sequence [8, 12]. As a consequence, the attribution of a video frame to a specific device can represent a much more challenging task than common image-camera attribution.

In this paper, we exploit PRNU-based traces to investigate the problem of device attribution when testing in-camera stabilized video frames. Specifically, given a device fingerprint \mathbf{K} and a frame \mathbf{I} coming from a stabilized video sequence, we aim at exploiting the PRNU traces left on \mathbf{I} in order to detect whether it has been recorded by the analyzed device. To do so, we assume that geometric transformation can be approximated by similarities [7, 8] and we propose

a geometrical realignment strategy based on a modified version of FM transform applied to both the device fingerprint and the frame noise residual. Specifically, the proposed modified Fourier-Mellin transform (MFM) enables comparing a device fingerprint and a noise residual independently from scaling and rotation operations. The next section provides all the details of the proposed method.

3. PROPOSED METHOD

In order to attribute a video frame \mathbf{I} to a device whose reference fingerprint is \mathbf{K} , we follow a pipeline based on a few steps: (i) noise extraction; (ii) geometric transformation estimation; (iii) geometric compensation and matching. In the following, we illustrate all the steps of the pipeline.

Noise extraction. As in common PRNU-based attribution algorithm, we extract the noise residual \mathbf{W} from frame \mathbf{I} . This is done using the common strategy proposed in [1, 11]: (i) noise is extracted through wavelet-based denoising; (ii) a series of post-processing steps (e.g., zero-averaging rows and columns, Wiener filtering, etc.) are applied to further enhance the noise residual \mathbf{W} .

Geometric transformation estimation. In order to match \mathbf{W} and the fingerprint \mathbf{K} , we first need to search for the geometrical transformation that might link them. In principle, assuming that video frames warping can be approximated by a similarity transformation [7, 8], aligning a noise residual and a reference device fingerprint by means of Fourier-Mellin may seem straightforward. In practice, differently from the Fourier-Mellin theory presented in Section 2, the two terms to compare (i.e., \mathbf{W} and \mathbf{K}) are not exactly one the transformed version (by means of a similarity transformation) of the other. First, the geometric transformation introduced by stabilization is not necessarily a similarity, but can include perspective distortions (on the entire frame or a localized portion of it) as well [12, 13]. Second, the noise residuals of video frames may contain scene content and noise contributions which are not present in the reference device fingerprint.

The primary consequence of this dissimilarity is that selecting only the Fourier magnitudes for estimating scale and rotation between the two terms, as reported in (3), may be not precise. Indeed, we verified that phase correlation between $\text{FM}\{\mathbf{K}\}$ and $\text{FM}\{\mathbf{W}\}$ does not show a pronounced peak, thus leading to a strongly hindered estimation of scale, rotation and shift. In order to overcome this issue we modify the Fourier-Mellin pipeline in two ways.

First, we propose to embed the phase term of the Fourier transform in addition to magnitude to the Fourier-Mellin pipeline. The modified Fourier-Mellin transform of \mathbf{I} can be thus defined as:

$$\text{MFM}\{\mathbf{I}\} = \text{LP}\{\mathbf{F}\}. \quad (4)$$

On one hand, phase adds more information, which is very useful for angle and scale estimation. On the other hand, this operation comes with a cost. The natural drawback of this approach is that we cannot isolate anymore the estimation of scale and rotation from the estimation of the shift. Indeed, in this case, phase correlation does not exclusively depends on scale and rotation transformations, but also on translation between the two terms. The Fourier-Mellin pipeline works only if \mathbf{W} and \mathbf{K} are almost perfectly aligned in terms of translation, i.e., if their mutual shift is basically 0 pixels in both horizontal and vertical directions. In other words, including the Fourier phase term, we first have to correctly realign the PRNU traces left on the noise residual with those on the reference fingerprint for what concerns the relative shift, then we can convert the Fourier transforms into log-polar domain and estimate the remaining parameters.




Fig. 1: Scheme of the proposed method for similarity estimation between noise residual \mathbf{W} and reference device fingerprint \mathbf{K} . The global optimizer searches for shift candidates, while the proposed MFM_{Δ_ρ} transform provides an estimate of scale and rotation for each shift.

The second proposed modification helps enabling faster computational times. It has been shown that a properly selected portion of the PRNU frequency spectrum can be sufficient to achieve good attribution performance (e.g., through subsampling [14]). In this vein, notice that a 2D frequency band becomes a rectangular band if the frequency spectrum is converted in log-polar domain. We propose to literally cut the frequency content of \mathbf{K} and \mathbf{W} by cropping the log-polar Fourier transform of Δ_ρ samples along the ρ dimension. The cropping center corresponds to the coordinate of the highest energy peak of $\text{MFM}\{\mathbf{K}\}$ evaluated as a function of ρ . Despite this step might seem irrelevant, this *strongly* reduces the amount of frequency samples to be correlated, thus lowering the computational cost. We define the modified Fourier-Mellin transform followed by cropping as:

$$\text{MFM}_{\Delta_\rho}\{\mathbf{I}\} = [\text{LP}\{\mathbf{F}\}]_{\Delta_\rho}. \quad (5)$$

By considering the added phase term and the frequency cropping step, the best similarity parameters can be estimated solving a maximization problem. Formally,

$$\hat{s}, \hat{\alpha}, \hat{c} = \arg \max_{s \in \mathcal{S}, \alpha \in \mathcal{A}, c \in \mathcal{C}} \Phi[\text{MFM}_{\Delta_\rho}\{\mathbf{W}\}, \text{MFM}_{\Delta_\rho}\{\mathbf{K}(\mathbf{x} - \mathbf{c})\}], \quad (6)$$

where Φ represents the phase correlation and vector \mathbf{x} refers to horizontal and vertical pixel coordinates.

Notice that, for each shift candidate value, scale and rotation parameters can be very quickly estimated in closed form through phase correlation. Therefore, we only need to optimize over different shift values. However, gradient descent strategies to solve (6) suffer from the non-convex behavior of the phase correlation as a function of the shift. Especially in video sequences characterized by outdoor scenarios or user motion, the actual peak value can be hard to find with gradient descent algorithms. The maximization problem as a function of the shift can be solved by resorting to global optimization techniques. It is worth noting that the translation between \mathbf{W} and \mathbf{K} can be assumed with slight approximation to imply integer shift in horizontal and vertical directions, i.e., to represent a certain number of pixels. We propose to exploit a global optimization algorithm known as genetic algorithm that allows an efficient estimation of integer parameters [15]. In a nutshell, the proposed method is shown in Fig. 1.

Geometric compensation and matching. After estimating the similarity transformation $\mathcal{T}_{\hat{s}\hat{\alpha}\hat{c}}$, last steps consist in: (i) applying $\mathcal{T}_{\hat{s}\hat{\alpha}\hat{c}}$ to \mathbf{K} in order to realign the PRNU traces left on \mathbf{K} with those of \mathbf{W} ; (ii) resorting to PCE as strategy for a correct source device identification. We compute P_{MFM} as

$$P_{\text{MFM}} = \text{PCE}(\mathbf{W}, \mathcal{T}_{\hat{s}\hat{\alpha}\hat{c}}(\mathbf{K})). \quad (7)$$

As in standard PRNU attribution tests, by thresholding P_{MFM} it is

possible to detect whether the frame under analysis belongs to the tested device. In case multiple frames are available, it is possible to repeat the whole procedure and fuse results obtained with different frames (e.g., maximum PCE picking, majority voting, etc.).

4. EXPERIMENTAL ANALYSIS

In this section we report all the details about the performed experimental campaign and the achieved results.

Dataset. Our datasets have been extracted from Vision dataset, which includes both images and videos from 11 major brand devices [9]. For building the PRNU related to each device, we select all the available images taken by the device depicting flat scenes [11]. Then, each fingerprint \mathbf{K} is built by scaling and cropping the PRNU, using the image to video warping parameters reported in [8]. Regarding video frames, we select only devices with Full-HD video resolution (i.e., 1920×1080 pixels). For the sake of clarity, we make use of the same device nomenclature presented in [9], creating two test datasets: a *non-stabilized* dataset, selecting non-stabilized devices D03, D11, D17, D21, D24, which come from 5 different brands, and a *stabilized* dataset, that includes all the 14 available stabilized devices.

Notice that the considered video frames contain both static and motion scenes, depicted as *still*, *panrot*, *move* in [9], and can include almost flat content as well as significant texture presence, denoted as *flat*, *indoor*, *outdoor* in [9]. In particular, we only make use of the I-frames, as the PRNU traces left on them are likely to be more reliable than inter-predicted frames [16, 6]. Furthermore, in light of past investigations about the first I-frame of stabilized video sequences, we always discard it from the experiments [8, 12].

MFM parameters. To compute the MFM transform, we evaluate the 2D Fourier transform over 4096×4096 samples after zero-padding residue and reference fingerprint in the pixel domain, in order not to introduce undesired border effects. Then, we convert both terms into log-polar domain, following the default parameters provided by [17], ending up with MFM transforms having 2896 ρ -samples and 2281 α -samples. We verified the chosen sampling grid for ρ and α dimensions allows a correct estimation of scaling factor and rotation angle. Eventually, we crop MFM transforms along ρ dimension according to the chosen number of samples Δ_ρ .

The exploited genetic algorithm mimics biological evolution to find a reliable shift estimation. Precisely, it has the following parameter configuration: a population size of 50 individuals, which iteratively update the cost function for a maximum of 50 iterations. Remaining parameters are those defined in [15].

Performance in a controlled scenario. In order to assess the accuracy in attributing video frames to the correct device, we investigate the proposed method in a controlled scenario. Specifically, considering the *non-stabilized* dataset, we randomly select 27 I-frames per device, taking care of equally distributing motion and static scenes, as well as flat and textured content. We end up with a total amount of 135 video frames. In particular, we select only frames which report acceptable PCE values with the device fingerprint (i.e., $\text{PCE} \geq 60$, as suggested in [18, 6, 8]). Then, we warp each frame by means of a similarity transformation, randomly selecting the parameters from some realistic ranges [12], namely $\mathcal{S} = [0.9, 1.1]$, $\mathcal{A} = [-3, 3]$ deg, $\mathcal{C} = [-90, 90]$ pixels, related to scale, rotation angle, horizontal and vertical shifts, respectively. We verified these ranges include the vast majority of possible similarity transformations between stabilized video frames and reference fin-

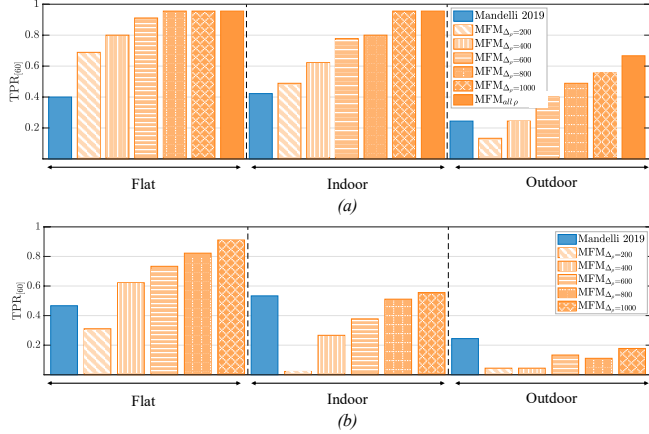


Fig. 2: Accuracy on synthetically warped non-stabilized video frames: (a) only scale and rotation are applied; (b) a complete similarity is applied. The proposed method (orange) can be tuned to use different amount of frequency samples ($\Delta\rho$), thus becoming slower but more accurate.

gerprint.

We aim at estimating the applied transformation using the proposed strategy, comparing the performance with the method presented in [8]. Specifically, we exploit the same parameter configuration for the particle swarm strategy [19, 15] as reported in [8], which enables to estimate the similarity transformation returning the highest PCE between \mathbf{W} and \mathbf{K} . For what regards the search bounds of scale and rotation parameters, we suppose these to be known at investigation side, thus they coincide with \mathcal{S} and \mathcal{A} . Notice that method [8] does not need to fix bounds for the shift parameters as these can be estimated without the need of optimization. Following similar considerations, the proposed MFM strategy fixes the search range for shift parameters exactly to \mathcal{C} , while scale and rotation do not require optimization.

Computational time and true positive attribution rate evaluated for a PCE threshold of 60 (i.e., $\text{TPR}_{[60]}$) are the chosen accuracy metrics to compare the two strategies. The average time for estimating the similarity transformation on each frame with the method [8] is 41s, while MFM strategy changes its temporal requirement depending on $\Delta\rho$ (e.g. using $\Delta\rho = 200$ requires only 21s on average). Generally, the required time linearly grows with $\Delta\rho$.

Fig. 2 shows results as a function of scene content of video frames (i.e., flat, indoor and outdoor). Specifically, Fig. 2(a) reports results where only scale-rotation transformations were applied. The shift between noise residuals and \mathbf{K} is assumed to be known. Fig. 2(b) reports results where a complete similarity transformation has been applied. It is worth noting that, in case the shift parameter is known and only scale and rotation parameters should be estimated, our proposal can be a viable solution for very fast identification. Since scale and rotation can be estimated without the need of optimization, the computational time reduces to less than one second. The more the selected $\Delta\rho$ samples, the better the accuracy of MFM strategy, which overcomes results of [8]. Furthermore, in this case there is no need for global optimizers, thus the potential optimization error reduces to zero. In case (b), MFM shows better or basically equivalent results to [8] for flat and indoor scenarios, while outdoor frames seems to be more challenging for the proposed method.

Performance on stabilized videos. In order to show the potentiality of MFM approach in dealing with source device identification problem on real videos, we apply the proposed method to the stabi-




Fig. 3: ROC curves obtained testing 10 I-frames with the proposed strategy, as a function of the number of used frequency samples, i.e., $\Delta\rho$. Results are compared to those of [8] evaluated using 10 I-frames.

Table 1: AUC and $\text{TPR}_{@0.01}$ testing 10 random I-frames per video query, together with average computational time per frame, evaluated with MFM and [8] methods.

Method	MFM $_{\Delta\rho=200}$	MFM $_{\Delta\rho=400}$	MFM $_{\Delta\rho=600}$	MFM $_{\Delta\rho=800}$	[8]
AUC	0.93	0.93	0.94	0.97	0.97
$\text{TPR}_{@0.01}$	0.85	0.86	0.90	0.94	0.91
Time [s]	22	46	69	105	57

lized video sequences. Following previous considerations, we set as search range for the mutual shift $\mathcal{C} = [-90, 90]$ both in horizontal and vertical directions. For clarity's sake, we use the very same accuracy metrics presented in [8], i.e., the area-under-the-curve AUC and $\text{TPR}_{@0.01}$ of ROC curves, averaged over all devices. Precisely, $\text{TPR}_{@0.01}$ corresponds to the rate of correct attributions evaluated when the false positive attribution rate is equal to 0.01.

We show the attribution results achieved by testing 10 random I-frames per video query and picking the maximum value among the computed PCEs. Specifically, we test different values for the number of used frequency samples (i.e., $\Delta\rho$) and always report results achieved by [8] over the same dataset. Fig. 3 draws the ROC curves and Table 1 depicts the achieved AUC and $\text{TPR}_{@0.01}$ as a function of $\Delta\rho$. Moreover, last row of Table 1 reports the average required computational time [seconds] for testing one query frame according to the chosen strategy, considering matching cases as well as non-matching ones. It is worth noticing that the proposed approach can overcome results of [8], provided that a sufficient amount of frequency samples is selected. Furthermore, the MFM strategy enables fast computations as well, at the expense of a slightly reduced accuracy, but still acceptable.

5. CONCLUSIONS

In this paper, we propose an alternative solution for solving the source device identification problem on stabilized videos. Specifically, we re-synchronize video frames and device reference fingerprint by estimating the re-alignment transformation with a modified version of the Fourier-Mellin transform. In doing so, we search the scaling and rotation parameters in the frequency domain, whereas unknown translation can be estimated leveraging global optimization strategies. Moreover, we propose to use a reduced amount of Fourier-Mellin transform samples to estimate the warping configuration, thus enabling fast computations.

The experimental campaign is conducted on a publicly available dataset. Results are promising and show enhanced performance with respect to state-of-the-art. This is especially true in situations where only scale and rotation parameters should be estimated: experiments performed in a synthetic set-up reveal that the proposed method can be much faster and accurate than existing methodologies.

6. REFERENCES

- [1] Jan Lukas, Jessica Fridrich, and Miroslav Goljan, "Digital camera identification from sensor pattern noise," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 205–214, 2006.
- [2] Nicola Mondaini, Roberto Caldelli, Alessandro Piva, Mauro Barni, and Vito Cappellini, "Detection of malevolent changes in digital video for forensic applications," in *Security, steganography, and watermarking of multimedia contents IX*. International Society for Optics and Photonics, 2007, vol. 6505, p. 65050T.
- [3] Aniello Castiglione, Giuseppe Cattaneo, Maurizio Cembalo, and Umberto Ferraro Petrillo, "Experimentations with source camera identification and online social networks," *Journal of Ambient Intelligence and Humanized Computing*, vol. 4, no. 2, pp. 265–274, 2013.
- [4] Flavio Bertini, Rajesh Sharma, Andrea Ianni, Danilo Montesi, and Mura Anteo Zamboni, "Social media investigations using shared photos," in *Proceedings of the International Conference on Computing Technology, Information Security and Risk Management (CTISRM)*, 2016, p. 47.
- [5] Miroslav Goljan and Jessica Fridrich, "Camera identification from cropped and scaled images," in *Security, Forensics, Steganography, and Watermarking of Multimedia Contents X*. International Society for Optics and Photonics, 2008, vol. 6819, p. 68190E.
- [6] Samet Taspinar, Manoranjan Mohanty, and Nasir Memon, "Source camera attribution using stabilized video," in *IEEE International Workshop on Information Forensics and Security (WIFS)*, 2016, pp. 1–6.
- [7] Massimo Iuliani, Marco Fontani, Dasara Shullani, and Alessandro Piva, "Hybrid reference-based video source identification," *Sensors*, vol. 19, no. 3, 2019.
- [8] Sara Mandelli, Paolo Bestagini, Luisa Verdoliva, and Stefano Tubaro, "Facing device attribution problem for stabilized video sequences," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 14–27, 2019.
- [9] Dasara Shullani, Marco Fontani, Massimo Iuliani, Omar Al Shaya, and Alessandro Piva, "VISION: a video and image dataset for source identification," *EURASIP Journal on Information Security*, vol. 2017, no. 1, pp. 15, Oct 2017.
- [10] B Srinivasa Reddy and Biswa Nath Chatterji, "An fft-based technique for translation, rotation, and scale-invariant image registration," *IEEE transactions on image processing*, vol. 5, no. 8, pp. 1266–1271, 1996.
- [11] Mo Chen, Jessica Fridrich, Miroslav Goljan, and Jan Lukáš, "Determining image origin and integrity using sensor noise," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 1, pp. 74–90, 2008.
- [12] Matthias Grundmann, Vivek Kwatra, and Irfan Essa, "Cascaded camera motion estimation, rolling shutter detection, and camera shake detection for video stabilization," Feb. 6 2018, US Patent 9,888,180.
- [13] Zhongqiang Wang, Lei Zhang, and Hua Huang, "High-quality real-time video stabilization using trajectory smoothing and mesh-based warping," *IEEE Access*, vol. 6, pp. 25157–25166, 2018.
- [14] Luca Bondi, Paolo Bestagini, Fernando Perez-Gonzalez, and Stefano Tubaro, "Improving PRNU compression through pre-processing, quantization and coding," *IEEE Transactions on Information Forensics and Security (TIFS)*, vol. 14, pp. 608–620, 2019.
- [15] Mathworks, "Global Optimization Toolbox - MATLAB R2016a," <https://www.mathworks.com/products/global-optimization.html>, 2016.
- [16] Wei-Hong Chuang, Hui Su, and Min Wu, "Exploring compression effects for improved source camera identification using strongly compressed video," in *IEEE International Conference on Image Processing (ICIP)*, 2011.
- [17] Mathworks, "Image Processing Toolbox- MATLAB R2016a," <https://www.mathworks.com/products/image.html>, 2016.
- [18] John Entrieri and Matthias Kirchner, "Patch-based desynchronization of digital camera sensor fingerprints," *Electronic Imaging*, vol. 2016, no. 8, pp. 1–9, 2016.
- [19] James Kennedy, "Particle swarm optimization," in *Encyclopedia of Machine Learning*, pp. 760–766. Springer, 2011.