

Secure Collaborative Camera Attribution*

A. Pedrouzo-Ulloa, F. Pérez-González, D. Vázquez-Padín

{apedrouzo, fperez, dvazquez}@gts.uvigo.es, atlantTic, Universidade de Vigo, Vigo, Galicia, Spain

ABSTRACT

Collaboration between law enforcement agencies (LEAs) is fundamental to successfully deal with the increasing amount of multimedia content in cross-border investigations. However, the fruitful exchange of sensitive data is highly limited due to privacy law restrictions. We propose a secure federated framework for camera attribution which enables LEAs to jointly exploit their forensic data.

CCS CONCEPTS

• **Security and privacy** → **Privacy-preserving protocols**; • **Applied computing** → **Evidence collection, storage and analysis**.

KEYWORDS

camera fingerprints, collaborative camera attribution, privacy protection, homomorphic encryption, multi-key encryption

ACM Reference Format:

A. Pedrouzo-Ulloa, F. Pérez-González, D. Vázquez-Padín. 2022. Secure Collaborative Camera Attribution. In *Proceedings of the European Interdisciplinary Cybersecurity Conference (EICC 2022)*, June 15–16, 2022, Barcelona, Spain. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3528580.3532993>

1 INTRODUCTION

The amount of multimedia content that law enforcement agencies (LEAs) must deal with in their investigations is ballooning. Moreover, collaboration between LEAs is becoming essential in a growing number of cases. However, the mere exchange of multimedia is strongly limited by privacy and data-protection laws, which may jeopardize a successful resolution.

Camera attribution problem. One especially relevant scenario within the field of image forensics is the case of camera attribution, for which the use of *camera device fingerprints* is very helpful to adequately determine the camera model and/or even the specific camera device which was used to take a picture. In particular, the Photo Response Non-Uniformity (PRNU) [10] is a visually imperceptible noise-like signal that is present in all images or videos taken by the same camera. Due to random imperfections produced during the manufacturing process, it acts as a unique identifier of a camera device. Since the first method for camera identification was proposed [6], this area of research has witnessed a considerable activity. Another breakthrough has been the inception of “*noiseprint*,”

a model-dependent fingerprint that is extracted using a Siamese DNN and may complement the use of the PRNU [1].

Privacy-sensitive fingerprints and multimedia content. Although camera attribution tools find immediate application in fighting against crimes involving depictions of minors, the faintness of these tools dictates that a number of images from the same device or camera model must be pooled together in order to extract fingerprints reliably. The highly-sensitive nature of the involved content impedes to easily work with this type of information and, even worse, privacy risks are aggravated after some recent works have shown that camera fingerprints estimates can leak a considerable amount of information from the images used for extraction [11].

1.1 Related Art and Objectives

Secure camera attribution. Recently, several works have proposed the use of secure computation techniques to address the privacy risks that are inherent to the estimation of camera device fingerprints. The three most relevant secure methods propose, respectively, a fully unattended solution based on the use of modern lattice-based cryptosystems [3], a combination of trusted hardware and Homomorphic Encryption (HE) [7], and the use of Shamir’s secret sharing [8] to alleviate the high computational costs, which requires to assume a non-colluding assumption for the external parties. While the last two methods assume that the extraction is performed inside the LEA’s own infrastructure and only cover the case of fingerprint detection, [3] performs the whole pipeline of extraction/detection in the encrypted domain. Even so, the methods from [3] introduce a high computational overhead, being impractical for the secure evaluation of the fingerprint extraction phase.

Current limitations. A common trait of all of the above methods is that they assume an outsourcing scenario, in which the secure execution of fingerprint extraction is not fully evaluated. Also, a collaboration among several institutions, where the required raw data is typically held separately, is not explicitly addressed. In this case sharing information is not always possible due to the existence of strict data protection and privacy regulations, such as the EU-promoted General Data Protection Regulation (GDPR), which strongly restricts the exchange of personal data across organizations in order to protect individuals’ rights. To the best of our knowledge, *a federated framework has never been proposed for the camera attribution problem.*

2 COLLABORATIVE FORENSIC SCENARIO

Collaboration among Data Owners. We consider a scenario in which the Data Owners are LEAs or Forensic Institutions aiming at camera fingerprint extraction through collectively trained models. The aggregator can be a Data Owner or, rather, a larger organization like Europol or Interpol. The end users would be the participating or other LEAs or Forensic institutions. As stated, the images used to extract the fingerprints may be very sensitive (e.g., in child abuse cases) and it has been shown that the high-pass residuals in the extraction may leak a considerable amount of information, so it

*This work was partially funded by the European Regional Development Fund (FEDER) and Xunta de Galicia under project “Grupos de Referencia Competitiva” (ED431C 2021/47). The first author is currently a visiting researcher at CEA-List, Université Paris-Saclay, being also funded by the European Union “NextGenerationEU/PRTR” by means of a Margarita Salas grant of the Universidade de Vigo.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

EICC 2022, June 15–16, 2022, Barcelona, Spain

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9603-5/22/06.

<https://doi.org/10.1145/3528580.3532993>

is desirable that those residuals, locally computed by each Data Owner, be stored and protected, even when used as training data for computing local fingerprint estimates and/or machine-learning (ML) model updates.

Privacy tools. To instantiate our framework we resort to multi-key HE (MKHE) [2, 5]. This tool finds a perfect fit here, as it incorporates a sort of access control structure into the ciphertexts, preventing other agents from decrypting any intermediate/final results without the permission of the other framework participants.

Example framework instantiations. The proposed framework not only focuses on the federated training of global ML forensic models, but will also cover a wide set of functionalities in which: 1) a Data Owner holds an image and searches for other Data Owners possessing images taken from the same camera; 2) a Data Owner holds an image and wants to determine whether any other Data Owner possesses a fingerprint estimate corresponding to the suspect image; 3) several Data Owners wish to improve their own fingerprint estimates by collaboratively aggregating them.

3 PROPOSED COLLABORATIVE FRAMEWORK

In our envisioned forensic scenario (see Section 2) with multiple Data Owners providing private data, the direct application of the state-of-the-art secure camera attribution solutions [3, 7, 8] is not entirely successful. All of them aim at securely outsourcing part of the camera attribution task to the cloud, but do not effectively address a secure collaboration among several institutions.

Our Federated Framework. Our federated framework (see Figure 1) builds upon two key concepts: 1) unprotected local data is isolated in different silos, and 2) all information leaving a silo is previously encrypted with MKHE. Ideally, MKHE provides on-the-fly encrypted operations between ciphertexts which are encrypted under different Data Owners' keys. However, MKHE does not scale as well as threshold HE when increasing the number of involved parties [5]. Taking this into account, here we will follow the blueprint of the threshold multiparty homomorphic (MHE) scheme from [5]:

- All involved Data Owners generate individual secret and a collective public keys (*Setup* step from [5, Protocol 6]).
- Each Data Owner encrypts its data to be outsourced under a collective public key (*Input* step from [5, Protocol 6]).
- One of the entities will be in charge of computing a particular functionality f (*Evaluation* step from [5, Protocol 6]).
- Finally, all involved parties collaboratively decrypt the outputs (*Output* step from [5, Protocol 6]).

The role of the functionality f depends on the concrete camera attribution task which is required by the Data Owners:

- *Training of ML models for noiseprint extraction.* The functionality f is the aggregation of local models. Federated training of Neural Networks with MHE is addressed in [9], where aggregated models are never decrypted at intermediate steps.
- *Aggregation of local PRNU estimates.* The functionality f is implemented with a homomorphic addition, followed by a division after decryption.
- *Fingerprint matching.* The functionality f corresponds with a set of homomorphic scalar products among encrypted fingerprint estimates.




Figure 1: A federated framework for camera attribution.

- *Residuals matching and/or fingerprint/residual matching.* The functionality f corresponds to a set of homomorphic scalar products with encrypted fingerprints/residuals.

The above functionalities can be implemented very efficiently with HE, as at worst they require arithmetic circuits with a multiplicative depth of 1. Also, this framework can be easily extended to encompass any other operation which is deemed adequate for a successful collaboration among LEAs and/or Forensic Institutions. Moreover, to enhance security, the execution of the local training/extraction in each Data Owners' own infrastructure can be performed by means of a trusted environment, similarly to [7].

Framework Example Instantiations. According to the single-key BFV ciphertext addition/multiplication runtimes¹ reported in [4], integer vectors of size 4096 can be added in less than 20 μ s and multiplied in less than 2100 μ s. Also, the number of involved parties (Data Owners) increases linearly the size of the underlying noise (see [5, Eq.(8)]). This means that even for 100 Data Owners, the size of the ciphertext modulo grows less than 8 bits, which does not increase significantly addition/multiplication runtimes for the MHE version of BFV with respect to single-key BFV.

Based on this assumption: 1) the encrypted aggregation of 100 local 2048×2048 PRNU fingerprint estimates would take around 2 seconds, while 2) the encrypted matching of a residual with 100 different PRNU fingerprint estimates (both of size 2048×2048) would take around 217 seconds. *These preliminary runtimes showcase the feasibility of a secure collaborative framework for camera attribution.*

REFERENCES

- [1] D. Cozzolino and L. Verdoliva. 2020. Noiseprint: A CNN-Based Camera Model Fingerprint. *IEEE Trans. Inf. Forensics Secur.* 15 (2020), 144–159.
- [2] A. López-Alt et al. 2017. Multikey Fully Homomorphic Encryption and Applications. *SIAM J. Comput.* 46, 6 (2017), 1827–1892.
- [3] A. Pedrouzo-Ulloa et al. 2018. Camera Attribution Forensic Analyzer in the Encrypted Domain. In *IEEE WIFS*. 1–7.
- [4] C. Mouchet et al. 2020. Lattigo: a Multiparty Homomorphic Encryption Library in Go. In *WAHC 2020*.
- [5] C. Mouchet et al. 2021. Multiparty Homomorphic Encryption from Ring-Learning-with-Errors. *Proc. Priv. Enhancing Technol.* 2021, 4 (2021), 291–311.
- [6] M. Chen et al. 2007. Digital imaging sensor identification (further study). In *Proc. SPIE*, Vol. 6505.
- [7] M. Mohanty et al. 2021. e-PRNU: Encrypted Domain PRNU-Based Camera Attribution for Preserving Privacy. *IEEE Tran. Dep. Sec. Comp.* 18, 1 (2021), 426–437.
- [8] R. Jena et al. 2021. SSS-PRNU: Privacy-Preserving PRNU Based Camera Attribution using Shamir Secret Sharing. *CoRR abs/2106.07029* (2021).
- [9] S. Sav et al. 2021. POSEIDON: Privacy-Preserving Federated Neural Network Learning. In *NDSS 2021*. The Internet Society.
- [10] Gerald C Holst. 1998. *CCD Arrays, Cameras, and Displays* (2nd ed.). SPIE Optical Engineering Press Bellingham, WA.
- [11] F. Pérez-González and S. Fernández-Mendiña. 2020. PRNU-leaks: facts and remedies. In *EUSIPCO*. IEEE, 720–724.

¹Runtimes were conducted single-threaded in [4] with the Lattigo library, so there is still room for significant parallelization.