

Compressed Fingerprint Matching and Camera Identification via Random Projections

Original

Compressed Fingerprint Matching and Camera Identification via Random Projections / Valsesia, Diego; Coluccia, Giulio; Bianchi, Tiziano; Magli, Enrico. - In: IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. - ISSN 1556-6013. - STAMPA. - 10:7(2015), pp. 1472-1485. [10.1109/TIFS.2015.2415461]

Availability:

This version is available at: 11583/2596759 since: 2022-08-22T14:10:01Z

Publisher:

IEEE - INST ELECTRICAL ELECTRONICS ENGINEERS INC

Published

DOI:10.1109/TIFS.2015.2415461

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Compressed Fingerprint Matching and Camera Identification via Random Projections

Diego Valsesia, *Student Member, IEEE*, Giulio Coluccia, *Member, IEEE*, Tiziano Bianchi, *Member, IEEE*, and Enrico Magli, *Senior Member, IEEE*

Abstract—Sensor imperfections in the form of photoresponse nonuniformity (PRNU) patterns are a well-established fingerprinting technique to link pictures to the camera sensors that acquired them. The noise-like characteristics of the PRNU pattern make it a difficult object to compress, thus hindering many interesting applications that would require storage of a large number of fingerprints or transmission over a bandwidth-limited channel for real-time camera matching. In this paper, we propose to use real-valued or binary random projections to effectively compress the fingerprints at a small cost in terms of matching accuracy. The performance of randomly projected fingerprints is analyzed from a theoretical standpoint and experimentally verified on databases of real photographs. Practical issues concerning the complexity of implementing random projections are also addressed using circulant matrices.

Index Terms—Random projections, PRNU, image forensics.

I. INTRODUCTION

IMAGING sensor imperfections can be considered as a unique fingerprint identifying a specific acquisition device, enabling various important forensic tasks, such as device identification, device linking, recovery of processing history, detection of digital forgeries [1]. The most common camera fingerprint is the photo-response nonuniformity (PRNU) of the digital imaging sensor [2]. The PRNU is due to slight variations in the properties of individual pixels, which produce a noise-like, yet deterministic pattern affecting every image taken by a sensor. Several works demonstrate that the PRNU is a robust fingerprint, usually surviving processing like lossy compression and image resizing [3], [4].

In the case of PRNU, the camera fingerprint is essentially a pattern with the same size as the imaging sensor. Due to the wide availability of sensors counting tens of millions of pixels, a realistic database of a few thousand sensors will require to store more than 10^{10} individual pixel values in uncompressed format. In addition, the complexity of looking for a particular fingerprint in a large database is also very high,

typically requiring the computation of a correlation with each fingerprint in the database. The issue of compression of PRNU patterns does not arise when the results of device identification have to be used as evidence in the court of law, because that case typically involves small databases and requires the highest accuracy. Instead, large scale problems, such as image classification, clustering or image retrieval problems based on camera identities, involve a huge number of PRNU patterns. Hence, these problems call for techniques to efficiently store and query such databases. Another problem with PRNU fingerprints is that the test image should be geometrically aligned with the fingerprint in the database. A possible solution is to provide several versions of the same fingerprint with different scale and/or cropping factors [5], however at the cost of managing an even larger database.

Recently, several authors [6] started to address the problems related with the management of a large database of camera fingerprints. In [7] and [8], the authors propose a so-called *fingerprint digest*, which works by keeping only a fixed number of the largest fingerprint values and their positions, so that the resulting database is independent of the sensor resolution. An improved search strategy based on fingerprint digest is proposed in [9] and [10]. Fingerprint digests can also be used to ease fingerprint registration in case of geometrically distorted images, as shown in [11]. An alternative solution is to represent sensor fingerprints in binary-quantized form [12]: even though the size of binary fingerprints scales with sensor resolution, binarization can considerably speed-up the fingerprint matching process.

In this paper, we propose a novel technique to reduce the size of camera fingerprints based on *random projections*. Our idea is motivated by the Johnson-Lindenstrauss (JL) lemma [13], stating that a small set of points in a high dimensional space can be embedded into a lower dimensional space approximately preserving the distances between the points, and by recent results showing that random linear projections can provide such embeddings with high probability [14]. In the case of PRNU fingerprints, it is easy to show that preserving the distance between two fingerprints is equivalent to preserving the angle between them. Since PRNU fingerprints of different sensors are known to be highly uncorrelated, and thus to form wide angles, we can expect that also the angles between compressed fingerprints obtained by random projections will be wide. As a consequence, in this paper we adapt the standard correlation detector [1] to solve fingerprint matching and camera identification problems in the compressed domain.

Manuscript received July 23, 2014; revised November 4, 2014, January 12, 2015, and March 17, 2015; accepted March 17, 2015. Date of publication March 23, 2015; date of current version June 2, 2015. This work was supported by the European Research Council through the European Community Seventh Framework Programme (FP7/2007-2013) under Grant 279848. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. H. Vicky Zhao.

The authors are with Department of Electronics and Telecommunications, Politecnico di Torino, Turin 10129, Italy (e-mail: diego.valsesia@polito.it; giulio.coluccia@polito.it; tiziano.bianchi@polito.it; enrico.magli@polito.it).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIFS.2015.2415461

As to practical issues, the complexity of randomly projecting a large fingerprint is greatly reduced by employing partial circulant matrices [15], which are known to be almost as good as fully random matrices. Moreover, inspired both by the work of [12] and by recent results in compressed sensing literature [16], we propose a binary version of the compressed fingerprint that further reduces storage and computational requirements.

The paper is organized as follows. In Section II, we provide notations and definitions and we briefly review forensic tasks based on PRNU and random projections. The proposed compressive PRNU forensic systems are described in Section III, while theoretical performance is analyzed in Section IV. Extensive numerical results on different datasets are presented and discussed in Section V. Finally, in Section VI we draw some conclusions.

II. BACKGROUND

A. Notation and Definitions

We denote (column-) vectors and matrices by lowercase and uppercase boldface characters, respectively. The ℓ -th element of column vector \mathbf{v} is v_ℓ . The i -th column of the matrix \mathbf{A} is \mathbf{a}_i .

The notation $\mathbf{A} \cdot \mathbf{B}$ denotes the elementwise product between matrices \mathbf{A} and \mathbf{B} , while \mathbf{A}/\mathbf{B} denotes elementwise division.

The notation $\langle \mathbf{a}, \mathbf{b} \rangle$ denotes the scalar product between vectors \mathbf{a} and \mathbf{b} , and $\|\mathbf{a}\|_2 = \sqrt{\langle \mathbf{a}, \mathbf{a} \rangle}$.

The notation $d_H(\mathbf{a}, \mathbf{b})$ denotes the Hamming distance between $\mathbf{a}, \mathbf{b} \in \{0, 1\}^m$, where $d_H(\mathbf{a}, \mathbf{b}) = \frac{1}{m} \sum_{i=1}^m a_i \oplus b_i$ and \oplus denotes the XOR operator.

The notation $\mathbf{a} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ means that the random vector \mathbf{a} is Gaussian distributed, its mean is $\boldsymbol{\mu}$, and its covariance matrix is $\boldsymbol{\Sigma}$.

B. PRNU Forensics

PRNU [1], [2] of imaging sensors is a property unique to each sensor array due to the different ability of each individual optical sensor to convert photons to electrons. This difference is mainly caused by impurities in silicon wafers and its effect is a noise pattern affecting every image taken by that specific sensor. Hence, the PRNU can be thought of as a spread-spectrum *fingerprint* of the sensor used to take a specific picture or a set of pictures. The PRNU is multiplicative, *i.e.*, if an imaging sensor is illuminated ideally with a uniform intensity \mathbf{i} ,¹ neglecting other sources of noise, the output of the sensor will be $\mathbf{o} = \mathbf{i} + \mathbf{i} \cdot \mathbf{k}$, where \mathbf{k} represents the matrix characterizing the PRNU values.

\mathbf{k} exhibits the following properties. It has the same pixel size as the sensor, and carries enough information to make it unique to each sensor. It is universal in the sense that every optical sensor exhibits PRNU. It is present in each picture taken by a sensor except from completely dark ones (due to its multiplicative nature). It is stable under different environmental conditions and is robust to several signal processing operations.

The PRNU characterizing one sensor can be extracted from a set of images (typically, 20 to 50 smooth images are enough). The procedure to extract the fingerprint \mathbf{k} of a sensor from a set of pictures depends on the model used to characterize the optical sensor. Denoting with \mathbf{i} the incident light intensity, the sensor output \mathbf{o} can be modelled as

$$\mathbf{o} = g^\gamma \cdot [(1 + \mathbf{k}) \cdot \mathbf{i} + \mathbf{e}]^\gamma + \mathbf{q}, \quad (1)$$

where g^γ is the gamma correction (g is different for each color channel and γ is usually close to 0.45), \mathbf{e} accounts for other noise sources internal to the sensor while \mathbf{q} models external noise (*e.g.* quantization). The goal is to extract \mathbf{k} , so, after keeping the first order term in the Taylor expansion of $[(1 + \mathbf{k}) \cdot \mathbf{i} + \mathbf{e}]^\gamma$, the output image can be factorized as

$$\mathbf{o} = \mathbf{o}^{\text{id}} + \mathbf{o}^{\text{id}} \cdot \mathbf{k} + \tilde{\mathbf{e}}, \quad (2)$$

where $\mathbf{o}^{\text{id}} = (\mathbf{g}\mathbf{i})^\gamma$ is the ideal sensor output, $\mathbf{o}^{\text{id}} \cdot \mathbf{k}$ is the PRNU term and $\tilde{\mathbf{e}} = \gamma \mathbf{o}^{\text{id}} \cdot \mathbf{e}/\mathbf{i} + \mathbf{q}$ collects other sources of noise. Assuming to be able to obtain through proper filtering a denoised version of \mathbf{o} , referred to as \mathbf{o}^{dn} , then this can be used as an approximation of the ideal sensor output and subtracted from each side of (2) to obtain the so-called *noise residual*, which can be modeled as:

$$\mathbf{w} = \mathbf{o} - \mathbf{o}^{\text{dn}} = \mathbf{o} \cdot \mathbf{k} + \tilde{\mathbf{q}}, \quad (3)$$

where $\tilde{\mathbf{q}}$ accounts for $\tilde{\mathbf{e}}$ and for the non-idealities of the model [1]. Suppose now that a certain number $C \geq 1$ of images is available. Considering the pixels of the noise term $\tilde{\mathbf{q}}$ as zero-mean Gaussian noise with variance σ^2 and independent from the signal $\mathbf{o} \cdot \mathbf{k}$, for each image ℓ , $\ell = 1, \dots, C$, it can be written

$$\mathbf{w}^{(\ell)}/\mathbf{o}^{(\ell)} = \mathbf{k} + \tilde{\mathbf{q}}/\mathbf{o}^{(\ell)}, \quad \text{where } \mathbf{w}^{(\ell)} = \mathbf{o}^{(\ell)} - \mathbf{o}^{(\ell)\text{dn}}. \quad (4)$$

Under the above assumptions, the log-likelihood of $\mathbf{w}^{(\ell)}/\mathbf{o}^{(\ell)}$ given \mathbf{k} satisfies

$$L(\mathbf{k}) = -\frac{C}{2} \sum_{\ell=1}^C \log \left(2\pi \sigma^2 / (\mathbf{o}^{(\ell)})^2 \right) \quad (5)$$

$$+ \sum_{\ell=1}^C \left(\mathbf{w}^{(\ell)}/\mathbf{o}^{(\ell)} - \mathbf{k} \right)^2 / \left(2\sigma^2 / (\mathbf{o}^{(\ell)})^2 \right) \quad (6)$$

from which the maximum likelihood estimate $\hat{\mathbf{k}}$ can be obtained as

$$\hat{\mathbf{k}} = \frac{\sum_{\ell=1}^C \left(\mathbf{w}^{(\ell)} \cdot \mathbf{o}^{(\ell)} \right)}{\sum_{\ell=1}^C (\mathbf{o}^{(\ell)})^2} \quad (7)$$

From the Cramer-Rao bound, the variance of the estimator can be estimated as

$$\sigma_{\hat{\mathbf{k}}}^2 = \sigma^2 / \sum_{\ell=1}^C (\mathbf{o}^{(\ell)})^2, \quad (8)$$

from which we can notice that good photos for fingerprint evaluation are photos with high luminance (but not saturated) and smooth content (which lowers σ^2). To improve further the quality of the estimation, artifacts shared among cameras of the same brand or model can be removed by subtracting row and

¹In this section, all vector quantities are vectorized versions of images.

column averages. In the case of color images, the estimation must be performed separately on each color channel, *i.e.*, we must obtain $\hat{\mathbf{k}}_R$, $\hat{\mathbf{k}}_G$ and $\hat{\mathbf{k}}_B$. After that, a “global” grayscale PRNU fingerprint will be obtained applying the usual RGB-to-gray conversion

$$\hat{\mathbf{k}} = 0.3\hat{\mathbf{k}}_R + 0.6\hat{\mathbf{k}}_G + 0.1\hat{\mathbf{k}}_B. \quad (9)$$

Several forensic tasks can be performed using the aforementioned model for camera sensors.

- The *device identification* problem [3] (also known in the biometrics field as *verification*) tests whether a given picture was taken by a specific device. An estimate of the fingerprint of the device has been extracted in advance from a set of training pictures and stored in a database. The noise residual or a single-image fingerprint estimate is extracted from the query image and correlated with the fingerprint in the database. The original detector presented in [4] correlates the noise residual of the query image with the database fingerprint modulated by the query image intensity, denoted as $\text{corr}(\mathbf{w}, \mathbf{o} \cdot \hat{\mathbf{k}})$.
- The *device linking* problem [17] is presented with two images and must determine whether they have been acquired by the same device. The noise residuals of the two photos are correlated, namely $\text{corr}(\mathbf{w}_1, \mathbf{w}_2)$. We will not discuss this usage case in the remainder of the paper.
- The *fingerprint matching* problem (also known in the biometrics field as *identification*) is presented with a database of fingerprint estimates and a set of pictures acquired by the same camera, which can be used to extract a fingerprint estimate. The goal is determine which device in the database (if present) has acquired the given pictures. Essentially, $\text{corr}(\hat{\mathbf{k}}, \hat{\mathbf{k}}_i)$ is calculated for all fingerprints, and if one fingerprint yields a correlation that is large enough, it is declared to be correct.

C. Random Projections

As will be explained in detail in Section III, PRNU databases can rapidly grow in size. For this reason, a method to “compress” them is required, with slight or ideally no information loss. One possible option is represented by *Random Projections* (RP), a low-complexity and yet powerful method for dimensionality reduction. The idea of RP is to project the original n -dimensional data to an m -dimensional subspace, with $m < n$, using a random matrix $\Phi \in \mathbb{R}^{m \times n}$. Hence, a collection of N n -dimensional data $\mathbf{D} \in \mathbb{R}^{n \times N}$ is reduced to an m -dimensional subspace $\mathbf{A} \in \mathbb{R}^{m \times N}$ by

$$\mathbf{A} = \Phi \mathbf{D}. \quad (10)$$

The key property behind RP is the Johnson–Lindenstrauss lemma [13], concerning low-distortion embeddings of points from high-dimensional into low-dimensional Euclidean space. The lemma states that a small set of points in a high-dimensional space can be embedded into a space of much lower dimension in such a way that distances between the points are nearly preserved.

Lemma 1 (Johnson–Lindenstrauss): Let $\varepsilon \in (0, 1)$. For every set \mathcal{Q} of $|\mathcal{Q}|$ points in \mathbb{R}^n , if m is a positive integer such

that $m = \mathcal{O}(\ln(|\mathcal{Q}|/\varepsilon^2))$, there exists a Lipschitz mapping $f: \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that

$$(1 - \varepsilon)\|\mathbf{u} - \mathbf{v}\|_2^2 \leq \|f(\mathbf{u}) - f(\mathbf{v})\|_2^2 \leq (1 + \varepsilon)\|\mathbf{u} - \mathbf{v}\|_2^2$$

for all $\mathbf{u}, \mathbf{v} \in \mathcal{Q}$.

It has been shown that f can be taken as a linear mapping represented by a random matrix $\Phi \in \mathbb{R}^{m \times n}$, whose entries are randomly drawn from certain probability distributions [14], like the Gaussian or Rademacher distributions.

The properties of RP are strictly related to the field of Compressed Sensing [18], [19], and in particular to the Restricted Isometry Property (RIP) of the sensing matrices [20]. In particular, in [20] it is shown that sensing matrices Φ whose elements follow the aforementioned distributions respect the RIP as well as the JL lemma. One can think of the RIP as a JL lemma specific for sparse vectors. In fact, a matrix $\Phi \in \mathbb{R}^{m \times n}$ is said to satisfy the RIP with constant δ_κ if there exists a constant δ_κ such that

$$(1 - \delta_\kappa)\|\mathbf{u}\|_2^2 \leq \|\Phi_\kappa \mathbf{u}\|_2^2 \leq (1 + \delta_\kappa)\|\mathbf{u}\|_2^2, \quad (11)$$

where Φ_κ is every possible $m \times \kappa$ submatrix obtained by keeping κ columns of Φ and $\kappa < n$.

The techniques presented in this paper bear some similarity with techniques used in Locality Sensitive Hashing (LSH) [21], [22]. Unlike standard hashing techniques, where the aim of the hashing function is to avoid collisions of hashes of different objects, LSH is a hashing technique for large databases using hashing functions whose aim is to maximize the probability of collision for objects close to each other rather than far apart. Then, the gap between the probability of collision of the hashes of similar objects and the probability of collision of the hashes of different objects is further amplified by concatenating several hashing functions. This allows one to perform, for example, a nearest-neighbour research in a large database using the hash of the query point retrieving elements stored in buckets containing that point. Several LSH families have been discovered in literature, each of them allowing a random choice of hashing functions. Among them, one, dubbed *arccos*, bears some similarity with 1-bit Compressed Sensing [16] and with the techniques explained in this paper. In words, the hashing function consists in the sign of the random projections, obtained with a sensing matrix with independent and identically distributed entries.

However, LSH is concerned with creating an efficient data structure to solve the approximate nearest-neighbor problem, so that one does not have to perform exhaustive search over the whole database. On the contrary, this paper addresses dimensionality reduction to create a compact representation for storage and computational complexity reduction of the matching operations, without the concern of the creation of a data structure to avoid exhaustive search. Indeed, since matching fingerprints typically do not show large correlation values, current results on LSH [23] demonstrate that it is hard to substantially improve over exhaustive search. Hence, while in this case random projections are not very effective at creating efficient data structures to avoid exhaustive search, nevertheless this paper shows that they are effective at reducing the dimensionality of the fingerprints. This allows to obtain

significant savings in storage space, and to speed up matching complexity, albeit requiring exhaustive search over the entire database, thanks to the smaller dimensionality and/or fast comparison between binary vectors, when only the sign of the RPs is kept.

III. COMPRESSIVE PRNU FORENSICS

This section describes how to apply compression techniques based on random projections to the forensic tasks presented in Sec. II-B.

A. Fingerprint Matching

Camera fingerprints obtained as PRNU patterns can be approximated as white Gaussian noise, a typical assumption considered in the literature to study the performance of matching systems. This has some important consequences: first, PRNU patterns cannot be compressed by standard methods (*e.g.*, JPEG compression) because they lack the redundancy that could be exploited to perform compression. Furthermore, fingerprints are very incoherent with each other. By incoherence, we mean that two fingerprints have very low correlation, or in other words, representing them as points in an n -dimensional space, the angle between any pair of fingerprints is wide and close to orthogonality. In the fingerprint matching problem, we construct a dictionary of fingerprints of N known cameras, which can be represented as a matrix $\mathbf{D} \in \mathbb{R}^{n \times N}$. The goal of the classic fingerprint matching problem is finding the column that is most similar to a test fingerprint $\hat{\mathbf{k}} \in \mathbb{R}^n$ that is presented to the system. To this purpose, one of the most used similarity criteria is the correlation coefficient. We will consider the sample reflective correlation, defined as follows:

$$\rho(\hat{\mathbf{k}}, \mathbf{d}_i) = \frac{\langle \hat{\mathbf{k}}, \mathbf{d}_i \rangle}{\|\hat{\mathbf{k}}\|_2 \|\mathbf{d}_i\|_2}, \quad i = 1, \dots, N \quad (12)$$

We propose to compress the database and test fingerprint representing them through a small number of random projections. This operation can be seen as the product times an $m \times n$ sensing matrix Φ :

$$\mathbf{A} = \Phi \mathbf{D} \quad (13)$$

$$\mathbf{y} = \Phi \hat{\mathbf{k}} \quad (14)$$

Random projections can effectively reduce the dimension of the space the fingerprints live in thanks to the fact that they approximately preserve the geometry of the point cloud composed of the fingerprints. Since random projections approximately preserve the angle between any two fingerprints and since this angle is wide thanks to their incoherent nature, we can expect a compressive system to exhibit robust performance, while dramatically reducing the problem size. The system has to store the compressed dictionary \mathbf{A} and a way to generate the compressed fingerprint whenever a test pattern is presented, using the same Φ (typically the seed of a pseudorandom number generator is stored).

The first system design challenge is the choice of the sensing matrix: the most studied sensing matrices are made of realizations of independent and identically distributed (i.i.d.)

Algorithm 1 Dictionary Creation

Require: \mathbf{D}, ϕ

Ensure: \mathbf{A}

```

for  $i = 1, \dots, N$  do
   $\mathbf{a}_i \leftarrow \text{IFFT} [\text{FFT}[\mathbf{d}_i] \cdot \text{FFT}[\phi]]$ 
   $\mathbf{a}_i \leftarrow$  first  $m$  entries of  $\mathbf{a}_i$ 
end for

```

Algorithm 2 Matching

Require: $\mathbf{A}, \hat{\mathbf{k}}, \phi$

```

 $\mathbf{y} \leftarrow \text{IFFT} [\text{FFT}[\hat{\mathbf{k}}] \cdot \text{FFT}[\phi]]$ 
 $\mathbf{y} \leftarrow$  first  $m$  entries of  $\mathbf{y}$ 

```

```

for  $i = 1, \dots, N$  do

```

```

  if  $\rho(\mathbf{y}, \mathbf{a}_i) > \tau$  then

```

```

    Declare a match

```

```

  end if

```

```

end for

```

Gaussian random variables. Although they can provide the best performance in terms of geometry preservation, Gaussian matrices present some drawbacks which make their use in large scale problems fairly complex. First, one needs to generate nm random numbers, which can take a significant amount of time when n is in the order of several millions. In practice one cannot typically store the whole matrix as this would require too much memory, so only the seed of a pseudorandom number generator is stored and every time the matrix is generated on-the-fly. Second, the full matrix by vector multiplication must be carried out for each of the columns of the dictionary. In order to avoid such problems we propose to use partial circulant matrices. Such matrices generate the first row ϕ at random (*e.g.*, with i.i.d. Gaussian variables), and all the other rows are just circularly shifted versions of the first row. It has been observed that circulant matrices perform almost as well as fully random Gaussian matrices, and proofs of the JL lemma [24] and of the RIP [25] are available for such matrices. Circulant matrices provide great advantages because only the first row must be generated at random, and because fast multiplication is available through the FFT. Thanks to the use of the FFT, the product $\Phi \mathbf{D}$ can be implemented with $\mathcal{O}(Nn \log n)$ operations instead of $\mathcal{O}(Nmn)$. The results presented in the following sections hold for circulant matrices with randomized column signs, since the proofs in [24] hold for this kind of matrices. In practice, randomizing the column signs of the sensing operator amounts to randomizing the signs of the signal and using the original operator. However, since our signals of interest are noise-like sequences, the randomization of the signs has no effect and it is possible to omit it. An example of a compressive system employing circulant sensing matrix is shown in Algs. 1 and 2.

Further compression can be achieved by quantizing the measurements, instead of keeping the floating point values. Jacques *et al.* [16] have shown in the field of 1-bit compressed sensing that random measurements with binary quantization implement an embedding that approximately preserves the

angle between signals. Since the preservation of the angles is the main interest for the matching problem, we will also consider the case of binary random measurements obtained as:

$$\mathbf{A} = \text{sign}(\Phi\mathbf{D}) \quad (15)$$

In the case of binary measurements the correlation coefficient is replaced by the Hamming distance as test metric.

$$d_H(\mathbf{y}, \mathbf{a}_i), \quad i = 1, \dots, N \quad (16)$$

In Sec. IV we discuss how the Hamming distance tends to be concentrated around $d_S(\hat{\mathbf{k}}, \mathbf{d}_i) = \pi^{-1} \arccos(\langle \hat{\mathbf{k}}, \mathbf{d}_i \rangle)$, being $\arccos(\langle \hat{\mathbf{k}}, \mathbf{d}_i \rangle)$ the angle between two uncompressed fingerprints. The higher the correlation between fingerprints, the narrower the angle between them. Hence, the angle between two matching fingerprints is typically narrower than the angle between non-matching fingerprints. This is reflected on the binary random projections, where the Hamming distance between matching fingerprints is typically smaller than that between non-matching fingerprints. Binary random projections allow to compress significantly, while the performance degradation is limited. As we will show in Sec. V, the degradation due to binarization is small but it allows to obtain a significant gain in terms of space. Moreover, computing the Hamming distance is a very fast and efficient operation. Binarization of the fingerprints was considered by Bayram *et al.* [12] as an effective method to reduce storage requirements. We go one step further by showing that binarization of random projections is effective as well, while further reducing the storage and computational requirements and providing additional flexibility by modulating the number of random measurements. Binarization of the fingerprints themselves can be seen as a special case of the presented framework, in which the sensing matrix is the identity.

B. Camera Identification

The camera identification problem is conceptually very similar to the fingerprint matching scenario. The main difference is that a single test image is available instead of a set of them. Chen *et al.* [4] showed that the optimal detector for this problem correlates the noise residual of the image with a modulated version of the fingerprint stored in the database, where the modulating term is the test image. Extending this detector to the compressed domain is not possible because of the elementwise product between test image and the fingerprint in the database. Instead, we investigate the performance of two simplified detectors that can be readily mapped to the compressed domain. The first simplified detector correlates the noise residual \mathbf{w} of the test image with the fingerprint stored in the database. Essentially this system eliminates the modulating effect of the test image, thus it will be sub-optimal unless the test image is a constant pattern. It is sufficient to apply the sensing matrix to both noise residual and fingerprint to translate this detector to the compressed domain.

$$\rho(\mathbf{w}, \mathbf{d}_i) \mapsto \rho(\Phi\mathbf{w}, \Phi\mathbf{d}_i) \quad (17)$$

The second simplified detector considers the use of a fingerprint estimate $\hat{\mathbf{k}}$ extracted from the single test image instead

of the noise residual. This is accomplished by means of the same procedure described in Sec.II-B, albeit with $C = 1$. The detector then correlates this test fingerprint estimate with the fingerprint stored in the dictionary.

$$\rho(\hat{\mathbf{k}}, \mathbf{d}_i) \mapsto \rho(\Phi\hat{\mathbf{k}}, \Phi\mathbf{d}_i) \quad (18)$$

C. Handling Sensors With Different Resolutions

When dealing with multiple camera sensors, it is commonly observed that they exhibit many different sizes, hence this variety of resolutions must be handled. In practice, one can resort to several solutions such as cropping a standard portion of fixed sized of every image or zero-padding the extracted fingerprints or noise residuals to a standard dimension. Zero-padding is typically the implied method when similarity is computed finding the maximum of the cross-correlation function (or the maximum peak-to-correlation energy [1]). The proposed compressive method projects elements of any dimension to the same m -dimensional space and then all computations are performed in this space. When two vectors $\mathbf{u} \in \mathbb{R}^{n_1}$ and $\mathbf{v} \in \mathbb{R}^{n_2}$, with $n_2 > n_1$ are measured as $\mathbf{y} = \Phi^{(u)}\mathbf{u}$ and $\mathbf{z} = \Phi^{(v)}\mathbf{v}$, the largest sensing matrix contains the smallest as a submatrix, namely $\Phi^{(v)} = [\Phi^{(u)} \ \Phi']$.

D. Detection Metrics

The matching problem is concerned with finding the column of the dictionary that best matches a test compressed pattern. The test compressed fingerprint undergoes a binary hypothesis test for each column of the compressed dictionary. The two hypotheses are defined as:

H₀: the compressed test fingerprint and the reference are not from the same camera

H₁: the compressed test fingerprint and the reference are from the same camera

We reject the null hypothesis whenever the correlation coefficient (or Hamming distance in the binary case) is above (respectively, below) a predefined threshold τ .

First, we define the following events, referring to a single instance of the hypothesis testing problem, *i.e.*, a single column of the dictionary. These are standard definitions, used for example in [7].

- **False Alarm**: the null hypothesis was incorrectly rejected.
- **Detection**: the null hypothesis was correctly rejected.

False alarm corresponds to the case in which the current column of the dictionary is the compressed fingerprint of a different camera with respect to the compressed fingerprint under test, but a match is incorrectly declared. On the other hand, *detection* occurs when the current column of the dictionary is the compressed fingerprint of the same camera as the compressed fingerprint under test, and a match is correctly declared.

Since previously-defined events are restricted to a single column of the dictionary, we also introduce global events considering the dictionary as a whole. These events are also defined in [12] and [26].

- **False Acceptance**: the null hypothesis was rejected for at least one wrong camera.

- **True Detection:** the null hypothesis was rejected only for the correct camera.

False acceptance corresponds to the case in which all the columns of the dictionary are tested, and at least one column containing the compressed fingerprint of a different camera with respect to the compressed fingerprint under test is declared as a match. On the other hand, *true detection* occurs when all the columns of the dictionary are tested, and a match is declared *only* for the column corresponding to the same camera of the compressed fingerprint under test.

IV. SYSTEM PERFORMANCE

In this section we provide some theoretical results concerning the performance of the proposed compressive system. In particular, we focus on the fingerprint matching problem. We provide a general framework to characterize the performance with arbitrary sensing matrices and with 1-bit quantization. In order to evaluate the performance, we consider the following model for the fingerprints. The system is presented a corrupted version of a fingerprint, namely

$$\hat{\mathbf{k}} = \mathbf{d}_i + \mathbf{z},$$

where \mathbf{z} is additive white Gaussian noise, *i.e.*, $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I}_n)$. Then, the compressive matching system with real-valued measurements applies (14) to compute the random projections $\mathbf{y} = \Phi \hat{\mathbf{k}}$ of the test fingerprint and compares them with each column of the compressed dictionary \mathbf{A} , namely $\mathbf{a}_j = \Phi \mathbf{d}_j$. On the other hand, in case of binary measurements the system applies (15) to compute the binary random projections $\mathbf{y} = \text{sign}(\Phi \hat{\mathbf{k}})$ of the test fingerprint and compares them with each column of the binarized compressed dictionary \mathbf{A} , namely $\mathbf{a}_j = \text{sign}(\Phi \mathbf{d}_j)$. We now formally define the probabilities of the events introduced in the previous section. In case of real-valued random projections, the probability of detection is

$$P_{D(i)} = \mathbb{P}(\rho(\mathbf{y}, \mathbf{a}_i) > \tau),$$

and the probability of false alarm is

$$P_{FA(i,j)} = \mathbb{P}(\rho(\mathbf{y}, \mathbf{a}_j) > \tau), \quad \text{with } i \neq j.$$

In case of binary measurements, the probabilities are, respectively,

$$\begin{aligned} P_{D(i)} &= \mathbb{P}(d_H(\mathbf{y}, \mathbf{a}_i) < \tau) \\ P_{FA(i,j)} &= \mathbb{P}(d_H(\mathbf{y}, \mathbf{a}_j) < \tau), \quad \text{with } i \neq j. \end{aligned}$$

Moreover, the probabilities of true detection and false acceptance are related to the probability of detection and false alarm by

$$\begin{aligned} P_{T(i)} &= P_{D(i)} \prod_{j \neq i} (1 - P_{FA(i,j)}) \\ P_{F(i)} &= 1 - \prod_{j \neq i} (1 - P_{FA(i,j)}). \end{aligned}$$

A. ε -Stable Correlation Embeddings

In order to characterize any compression matrix we propose to introduce a property, dubbed *ε -stable correlation embedding*, that, if satisfied, allows to write bounds for the probability of false alarm and the probability of detection.

Definition 2: An ε -stable correlation embedding (ε -SCE) of a set \mathcal{P} of N points of \mathbb{R}^n is a map $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that:

$$\langle \mathbf{u}, \mathbf{v} \rangle - \varepsilon \leq \langle \varphi(\mathbf{u}), \varphi(\mathbf{v}) \rangle \leq \langle \mathbf{u}, \mathbf{v} \rangle + \varepsilon$$

for all $\mathbf{u}, \mathbf{v} \in \mathcal{P}$.

Essentially, we are requiring an approximate preservation of inner products when the compression matrix is applied to the fingerprints. Note that when the fingerprints are normalized to have unit norm ($\|\mathbf{d}_i\|_2 = 1$), the correlation coefficient corresponds to the inner product. If the sensing matrix is an ε -stable correlation embedding then it is easy to derive bounds on the false alarm and detection probabilities, in terms of the respective probabilities in the uncompressed case, as explained in the following.

Theorem 3: Let $\Phi \in \mathbb{R}^{m \times n}$ be an ε -SCE for a set of N camera fingerprints $\mathcal{P} = \{\mathbf{d}_i \in \mathbb{R}^n : \|\mathbf{d}_i\|_2 = 1, i = 1, \dots, N\}$, with $\varepsilon \in [0, 1)$, and $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma_z^2 \mathbf{I}_n)$. Then, the following bounds on the detection and false alarm probabilities hold:

$$\begin{aligned} P_{D(i)} &\geq P_{D(i)}^{\text{UN}} \left(\frac{\tau}{1 - \varepsilon} \right) \\ P_{FA(i,j)} &\leq P_{FA(i,j)}^{\text{UN}} \left(\frac{\tau - \varepsilon(1 - \langle \mathbf{d}_i, \mathbf{d}_j \rangle)}{1 + \varepsilon} \right), \end{aligned}$$

where

$$P_{D(i)}^{\text{UN}}(\tau) = \mathbb{P}(\langle \mathbf{d}_i + \mathbf{z}, \mathbf{d}_i \rangle > \tau) = \mathbb{P}(\langle \mathbf{z}, \mathbf{d}_i \rangle > \tau - 1)$$

and

$$\begin{aligned} P_{FA(i,j)}^{\text{UN}}(\tau) &= \mathbb{P}(\langle \mathbf{d}_i + \mathbf{z}, \mathbf{d}_j \rangle > \tau) \\ &= \mathbb{P}(\langle \mathbf{z}, \mathbf{d}_j \rangle > \tau - \langle \mathbf{d}_i, \mathbf{d}_j \rangle) \end{aligned}$$

are the probabilities of detection and false alarm, respectively, in the uncompressed domain.

Proof:

$$\begin{aligned} P_{D(i)} &= \mathbb{P}(\langle \Phi(\mathbf{d}_i + \mathbf{z}), \Phi \mathbf{d}_i \rangle > \tau) \\ &= \mathbb{P}(\langle \Phi \mathbf{d}_i, \Phi \mathbf{d}_i \rangle + \langle \Phi \mathbf{z}, \Phi \mathbf{d}_i \rangle > \tau) \\ &= \mathbb{P}(\langle \Phi \mathbf{z}, \Phi \mathbf{d}_i \rangle > \tau - \langle \Phi \mathbf{d}_i, \Phi \mathbf{d}_i \rangle) \\ &\geq \mathbb{P}(\langle \Phi \mathbf{z}, \Phi \mathbf{d}_i \rangle > \tau + \varepsilon - 1) \\ &= \mathbb{P}\left(\frac{\langle \Phi \mathbf{z}, \Phi \mathbf{d}_i \rangle}{\|\Phi^\top \Phi \mathbf{d}_i\|_2} > \frac{\tau + \varepsilon - 1}{\|\Phi^\top \Phi \mathbf{d}_i\|_2}\right) \end{aligned}$$

We know that $\frac{\langle \Phi \mathbf{z}, \Phi \mathbf{d}_i \rangle}{\|\Phi^\top \Phi \mathbf{d}_i\|_2}$ has the same distribution as $\langle \mathbf{z}, \mathbf{d}_i \rangle$ and that

$$(1 - \varepsilon) \leq \|\Phi^\top \Phi \mathbf{d}_i\|_2 \leq (1 + \varepsilon)$$

since, as a consequence of Def. 2, the minimum and maximum eigenvalues of $\Phi^T \Phi$ are $1 - \varepsilon$ and $1 + \varepsilon$, respectively. Hence,

$$\begin{aligned} P_{D(i)} &\geq \mathbb{P} \left(\langle \mathbf{z}, \mathbf{d}_i \rangle > \frac{\tau + \varepsilon - 1}{\|\Phi^T \Phi \mathbf{d}_i\|_2} \right) \\ &\geq \mathbb{P} \left(\langle \mathbf{z}, \mathbf{d}_i \rangle > \frac{\tau + \varepsilon - 1}{1 - \varepsilon} \right) = P_{D(i)}^{\text{UN}} \left(\frac{\tau}{1 - \varepsilon} \right). \end{aligned}$$

Correspondingly,

$$\begin{aligned} P_{FA(i,j)} &= \mathbb{P} (\langle \Phi(\mathbf{d}_i + \mathbf{z}), \Phi \mathbf{d}_j \rangle > \tau) \\ &= \mathbb{P} (\langle \Phi \mathbf{d}_i, \Phi \mathbf{d}_j \rangle + \langle \Phi \mathbf{z}, \Phi \mathbf{d}_j \rangle > \tau) \\ &= \mathbb{P} (\langle \Phi \mathbf{z}, \Phi \mathbf{d}_j \rangle > \tau - \langle \Phi \mathbf{d}_i, \Phi \mathbf{d}_j \rangle) \\ &\leq \mathbb{P} (\langle \Phi \mathbf{z}, \Phi \mathbf{d}_j \rangle > \tau - \varepsilon - \langle \mathbf{d}_i, \mathbf{d}_j \rangle) \\ &= \mathbb{P} \left(\frac{\langle \Phi \mathbf{z}, \Phi \mathbf{d}_j \rangle}{\|\Phi^T \Phi \mathbf{d}_j\|_2} > \frac{\tau - \varepsilon - \langle \mathbf{d}_i, \mathbf{d}_j \rangle}{\|\Phi^T \Phi \mathbf{d}_j\|_2} \right) \\ &= \mathbb{P} \left(\langle \mathbf{z}, \mathbf{d}_j \rangle > \frac{\tau - \varepsilon - \langle \mathbf{d}_i, \mathbf{d}_j \rangle}{\|\Phi^T \Phi \mathbf{d}_j\|_2} \right) \\ &\leq \mathbb{P} \left(\langle \mathbf{z}, \mathbf{d}_j \rangle > \frac{\tau - \varepsilon - \langle \mathbf{d}_i, \mathbf{d}_j \rangle}{1 + \varepsilon} \right) \\ &= P_{FA(i,j)}^{\text{UN}} \left(\frac{\tau - \varepsilon(1 - \langle \mathbf{d}_i, \mathbf{d}_j \rangle)}{1 + \varepsilon} \right). \end{aligned}$$

We can notice that the performance of the compressed system can be linked to the performance of the uncompressed system with a modified threshold. The detection and false alarm probabilities have a threshold that is respectively increased or decreased by a function of ε . For a better subspace embedding, ε tends to zero and the corresponding threshold approaches τ , thus the performance approaches the one of the uncompressed system. Closed form expressions for the probabilities in the uncompressed case are available in the literature [1].

We now prove the ε -SCE for some important sensing matrices. It can be remarked that a matrix satisfying the JL lemma, also represents an ε -SCE, since the preservation of inner products comes as a corollary to the preservation of the Euclidean norm. Thanks to the vast literature on JL embeddings, the results for many matrices of interest are readily available. We now adapt a few of them to the ε -SCE formulation. In particular, we consider the Gaussian case again because it is a classic result and the circulant case because of its practical use.

Theorem 4 (Gaussian Matrices): Given a set \mathcal{P} of N unit-norm points in \mathbb{R}^n , fix $\varepsilon > 0$, and let Φ be a random matrix whose entries are i.i.d. Normal random variables with zero mean and $\frac{1}{m}$ variance. Then, Φ is an ε -SCE of \mathcal{P} with probability exceeding $1 - N^2 e^{(2 - (\varepsilon^2 - \varepsilon^3)) \frac{m}{4}}$.

Proof: Applying the JL lemma for Gaussian matrices [27] to vectors $\mathbf{u} + \mathbf{v}$ and $\mathbf{u} - \mathbf{v}$ we know that with probability

exceeding $1 - 4e^{-(\varepsilon^2 - \varepsilon^3)m/4}$:

$$\begin{aligned} (1 - \varepsilon)\|\mathbf{u} + \mathbf{v}\|_2^2 &\leq \|\Phi(\mathbf{u} + \mathbf{v})\|_2^2 \leq (1 + \varepsilon)\|\mathbf{u} + \mathbf{v}\|_2^2 \\ (1 - \varepsilon)\|\mathbf{u} - \mathbf{v}\|_2^2 &\leq \|\Phi(\mathbf{u} - \mathbf{v})\|_2^2 \leq (1 + \varepsilon)\|\mathbf{u} - \mathbf{v}\|_2^2 \end{aligned}$$

We can rewrite the inner product as:

$$\begin{aligned} 4\langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle &= \|\Phi(\mathbf{u} + \mathbf{v})\|_2^2 - \|\Phi(\mathbf{u} - \mathbf{v})\|_2^2 \\ &\geq (1 - \varepsilon)\|\mathbf{u} + \mathbf{v}\|_2^2 - (1 + \varepsilon)\|\mathbf{u} - \mathbf{v}\|_2^2 \\ &= 4\langle \mathbf{u}, \mathbf{v} \rangle - 2\varepsilon \left(\|\mathbf{u}\|_2^2 + \|\mathbf{v}\|_2^2 \right) \geq 4\langle \mathbf{u}, \mathbf{v} \rangle - 4\varepsilon. \end{aligned}$$

Hence, we can write:

$$\mathbb{P}(|\langle \mathbf{u}, \mathbf{v} \rangle - \langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle| \geq \varepsilon) \leq 4e^{-(\varepsilon^2 - \varepsilon^3) \frac{m}{4}}.$$

Given a set \mathcal{P} of N points, there are $\binom{N}{2}$ possible $\langle \mathbf{u}, \mathbf{v} \rangle$. We can use the union bound to find a concentration of measure for any $\langle \mathbf{u}, \mathbf{v} \rangle$.

$$\begin{aligned} \mathbb{P}(|\langle \mathbf{u}, \mathbf{v} \rangle - \langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle| \geq \varepsilon) &\leq \binom{N}{2} 4e^{-(\varepsilon^2 - \varepsilon^3) \frac{m}{4}} \\ &\leq N^2 e^{(2 - (\varepsilon^2 - \varepsilon^3)) \frac{m}{4}}. \end{aligned}$$

Theorem 5 (Circulant Matrices): Given a set \mathcal{P} of N points in \mathbb{R}^n within the unit ball, fix $\varepsilon \in (0, \frac{1}{2})$, and let $\Phi = \Phi_c \mathbf{S}$ where Φ_c is an $m \times n$ circulant matrix with the first row being an i.i.d. sequence of Normal random variables with zero mean and $\frac{1}{m}$ variance or an i.i.d. sequence of Rademacher random variables rescaled by a factor $\frac{1}{\sqrt{m}}$, and \mathbf{S} is a diagonal matrix of i.i.d. equiprobable ± 1 . Then, Φ is an ε -SCE of \mathcal{P} with probability exceeding $1 - N^2 e^{(2 - c m \varepsilon^2) \frac{1}{3}}$, for some constant c .

Proof: The proof basically follows the one for Gaussian matrices and uses the result in [24], which shows that for any vector \mathbf{x} , the following holds:

$$\mathbb{P} \left((1 - \varepsilon)\|\mathbf{x}\|_2^2 \leq \|\Phi \mathbf{x}\|_2^2 \leq (1 + \varepsilon)\|\mathbf{x}\|_2^2 \right) \geq 1 - 2e^{-c(m\varepsilon^2)^{1/3}}$$

Applying it to $\mathbf{u} + \mathbf{v}$ and $\mathbf{u} - \mathbf{v}$ for a fixed pair $\mathbf{u}, \mathbf{v} \in \mathcal{P}$ we can derive:

$$\mathbb{P}(|\langle \mathbf{u}, \mathbf{v} \rangle - \langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle| \leq \varepsilon) \geq 1 - 4e^{-c(m\varepsilon^2)^{1/3}}$$

Finally, we can apply the union bound to all the points in \mathcal{P} , to obtain:

$$\mathbb{P}(|\langle \mathbf{u}, \mathbf{v} \rangle - \langle \Phi \mathbf{u}, \Phi \mathbf{v} \rangle| \leq \varepsilon) \geq 1 - N^2 e^{2 - c(m\varepsilon^2)^{1/3}}$$

We remark that according to the previous results, Gaussian matrices require $m = \mathcal{O}(\varepsilon^{-2} \log N)$ measurements, while circulant matrices require $m = \mathcal{O}(\varepsilon^{-2} \log^3 N)$ measurements. However, the latter result is not sharp and recent works [15], [28] focused on improving the bound, closing some of the gap with respect to the result on Gaussian matrices. This highlights the slight reduction on performance due to the circulant structure with respect to a fully random matrix.

B. Binary Measurements

We now analyze the case of binary measurements. The role previously played by the inner product is now played by the angle between uncompressed fingerprints \mathbf{u} and \mathbf{v} and the Hamming distance between their binary measurement vectors. Here we introduce a binary equivalent of Definition 2.

Definition 6: An ε -stable binary embedding (ε -SBE) of a set \mathcal{P} of N points of \mathbb{R}^n is a map $\varphi : \mathbb{R}^n \rightarrow \{0, 1\}^m$ such that:

$$d_S(\mathbf{u}, \mathbf{v}) - \varepsilon \leq d_H(\varphi(\mathbf{u}), \varphi(\mathbf{v})) \leq d_S(\mathbf{u}, \mathbf{v}) + \varepsilon \quad (19)$$

for all $\mathbf{u}, \mathbf{v} \in \mathcal{P}$.

In the binary case, it is difficult to rigorously derive closed-form bounds to the probabilities of detection and false alarm, analogue to Theorem 3. Nevertheless, we conjecture that the performance of the compressed binary system is related to the constant ε of the embedding ε -SBE, and tends to the performance of the uncompressed system as ε tends to zero. In our case, $\varphi(\mathbf{u}) = \text{sign}(\Phi\mathbf{u})$, and the conjecture is supported by experimental evidence, shown in Section V. In [16], the authors defined a property (called $B\varepsilon$ SE), which is similar to Definition 6, but restricted to κ -sparse signals. When Φ is a Gaussian matrix, [16] also showed that $\text{sign}(\Phi\mathbf{u})$ provides a $B\varepsilon$ SE with high probability. Since we are not concerned with covering the set of all κ -sparse signals, we can exploit [16, Lemma 2] followed by a union bound argument, as we did in the proof of Theorem 4, to prove that $\text{sign}(\Phi\mathbf{u})$ also provides an ε -SBE.

Extending the above theoretical result to circulant matrices is an open problem, but experimental results seem to confirm the validity of the ε -SBE.

V. EXPERIMENTAL RESULTS

We tested the performance of the compressed system under various conditions. We used two datasets of actual photographs to obtain the receiver operating characteristic (ROC) of the system under different scenarios. We constructed the first dataset (PoliTO database) by shooting photographs of walls with 8 different cameras. The uniform subject and the control over light conditions make those photos nearly ideal for the extraction of camera fingerprints. The second database is the publicly available Dresden image database [29]. Each database is constructed from a number of training photos, while T additional photos are used for testing. Extraction of the camera fingerprints is performed using the Camera Fingerprint toolbox [30], [31].

Referring to the events described in Sec. III-D and the probabilities defined in Section IV, we estimate the detection probability $P_{D(i)}$, averaged over all the cameras $i = 1, \dots, N$, with the *true positive rate* as

$$\text{True Positive Rate} = \frac{\# \text{ of detections}}{T},$$

while the false alarm probability $P_{FA(i,j)}$, averaged over all the cameras $i = 1, \dots, N$ and $j \neq i$, is estimated with the *false positive rate* as

$$\text{False Positive Rate} = \frac{\# \text{ of false alarms}}{(N-1)T}.$$

A first ROC plots the True Positive Rate vs. the False Positive Rate. The ideal curve is one for any False Positive Rate.

Next, we estimate the true detection probability $P_{T(i)}$, averaged over all the cameras $i = 1, \dots, N$, with the *true detection rate*, as

$$\text{True Detection Rate} = \frac{\# \text{ of true detections}}{T},$$

while the false acceptance probability $P_{F(i)}$, averaged over all the cameras $i = 1, \dots, N$, is estimated with the *false acceptance rate* as

$$\text{False Acceptance Rate} = \frac{\# \text{ of false acceptances}}{T}.$$

A second ROC plots the True Detection Rate vs. the False Acceptance Rate. The ideal curve is the top-left–bottom-right diagonal.

All the tests use detector (17), for reasons that will be explained in Section V-E, and compressed fingerprints are obtained using a circulant sensing matrix, whose performance is compared to the one obtained using Gaussian sensing matrices in Section V-F. We show the results of the following experiments. First, we show in Sections V-A, V-B, V-C and V-D how the choice of the dimension m affects the performance of the system, both for real-valued and binary random projections. It is clear that dimensionality reduction degrades the performance with respect to the uncompressed system, but we show that a suitable choice of m allows to significantly reduce the storage requirement and computational complexity with comparable performance. Moreover, Sections V-C and V-D report a comparison with other state-of-the-art methods for fingerprint compression, namely a fingerprint digest [9], [10] and fingerprint binarization [12].

A. Concentration of Correlation

In this section, we use simulated fingerprints to graphically show how the correlation between compressed fingerprints changes as a function of the number of measurements, consequently affecting the performance. This is a more intuitive representation of the theorems presented in Sec. IV. The synthetic fingerprints are generated as vectors of i.i.d. zero mean random Gaussian variables with unit variance. Fig. 1 shows the empirical distribution of correlation of matching and non-matching fingerprints. This figure has been created by generating 50000 synthetic fingerprints matching a reference fingerprint with correlation coefficient 0.05 ± 10^{-4} , and 50000 non-matching synthetic fingerprints (correlation: 0 ± 10^{-3}). First of all, we notice that the correlation between compressed fingerprints tends to concentrate around the original correlation between the uncompressed versions. Notice that the width of the concentration peak is determined by the number of measurements but does not depend on the original fingerprint size. This clearly appears in the Johnson-Lindenstrauss lemma which claims a dependency on m and N alone. It is evident that the original correlation plays a key role in the number of measurements to be selected because a sharper peak is needed for low correlations values.




Fig. 1. Empirical distribution of correlation between matching and non-matching compressed fingerprints. Uncompressed correlation coefficient for matching: $\rho = 0.05 \pm 10^{-4}$. Uncompressed correlation coefficient for non-matching: $\rho = 0 \pm 10^{-3}$.

TABLE I
SIZE OF POLITO DATABASE IN BYTES (8 CAMERAS)

	Single precision	Binary
$m = 1000$	31.25 KiB	1000 B
$m = 2000$	62.50 KiB	2000 B
$m = 4000$	125.00 KiB	4000 B
$m = 8000$	250.00 KiB	8000 B
Uncompressed	331.39 MiB	10.36 MiB

B. PoliTO Database

The PoliTO database is composed of pictures from 8 different consumer cameras. The pictures are defocused photos of walls under good illumination conditions. Each camera has at least 100 photos, all in landscape format, shot at the full resolution and maximum quality JPEG compression. We use 60 photos of each camera to extract the ground truth fingerprint to be stored in the database, while the remaining ones are used for testing purposes. Each ROC curve is obtained by sweeping the threshold parameter τ . Test images are presented to the system one at a time, the noise residual is extracted and then compressed using the same sensing matrix used to compress the database. Figure 2 shows various ROCs parametrized by the number of measurements. It can be noticed that a very small number of random measurements is enough to get almost indistinguishable performance from a perfect detector, while saving a considerable amount of storage space. Table I shows some actual figures for the space needed to store the dictionary of fingerprints on disk (without any additional form of lossless compression, which is anyway highly ineffective due to the high entropy of the PRNU and of the random measurements).

C. Dresden Database

The database assembled in [29] is composed of both flatfield images and scenes from indoor and outdoor environments. We selected 53 cameras having both flatfield and natural photos. The database is created from the flatfield images in order to have high quality fingerprints, while the test images are taken from the natural scenes. The natural photos

TABLE II
SIZE OF DRESDEN DATABASE IN BYTES (53 CAMERAS)

	Single precision	Binary
$m = 16000$	3.23 MiB	103.51 KiB
$m = 32000$	6.47 MiB	207.03 KiB
$m = 64000$	12.94 MiB	414.06 KiB
$m = 128000$	25.88 MiB	828.13 KiB
$m = 512000$	103.52 MiB	3.23 MiB
Uncompressed	1882.15 MiB	58.82 MiB

present varying amounts of details and illumination conditions, thus making this dataset much more challenging than the PoliTO database. All photos are registered to the same sensor orientation. Figure 3 shows the ROC curves parametrized by number of measurements. It is observed that some test photos are very challenging to match, thus the ceiling in the ROC. This has been previously observed in [32] where the authors observe that some camera models present uncommon non-unique artefacts, mostly due to on-board post-processing. As expected, a higher number of measurements is required for this database due to the lower correlation between the noise residual extracted from a test image and the fingerprint stored in the database. We remark that a lower correlation on the original uncompressed fingerprints implies a narrower angle between them, hence a higher number of random projections is required to preserve it with good accuracy. Table II shows some actual figures for the space needed to store the dictionary on disk. We also make a comparison with alternative compression techniques based on trimming the fingerprint to a fixed length [26] or creating a digest. The sequential trimming technique simply retains a fixed number of entries of the fingerprint in fixed locations, and then optionally quantizes the entries. This is equivalent to using a partial identity as a sensing matrix. However, a partial identity provides an embedding which is less robust against bad inputs such as the case of a localized strong noise in the retained area. The digest technique, proposed in [11], is an adaptive compression method that retains the d entries of the fingerprint with largest magnitude. We note that the complexity of the method is largely similar to the complexity of random projections; in fact the digest creation requires a $\mathcal{O}(n \log n)$ step to identify the d largest elements, while the proposed techniques require computing the random measurements, which is done in $\mathcal{O}(n \log n)$ time thanks to the FFT. However, the digest method requires to store the positions corresponding to the retained entries. The comparison presented in Fig. 4 shows two set of curves obtained for a fixed bit budget of 512000 bits and of 128000 bits. Binary-quantized random projections are compared with real-valued and binary-quantized fingerprints trims (d entries in the top left corner are retained) and real-valued and binary-quantized digests. The real-valued digest uses 32 bits for pixel intensity and 24 bits for location information, while the binary-valued digest uses only 1 bit for pixel intensity and 24 bits for location information. It can be noticed that binary random projections outperform all the other methods, that the digest is better than trimming and that binary quantization of the digest pixels marginally improves the performance. Finally, Fig. 5 reports a simulation on a synthetic database of $N = 10$ cameras ($n = 10 \cdot 10^6$) and

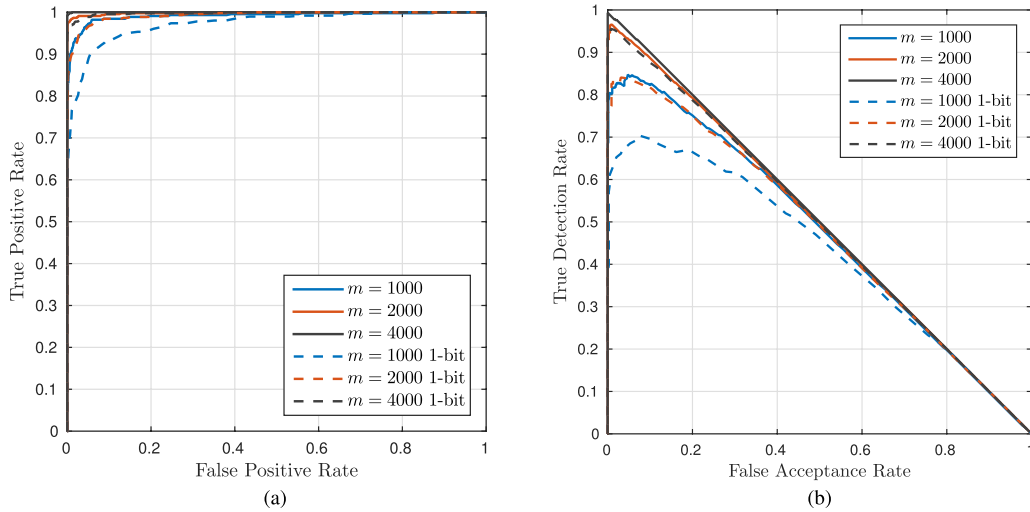


Fig. 2. ROC curves for the PoliTo database. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

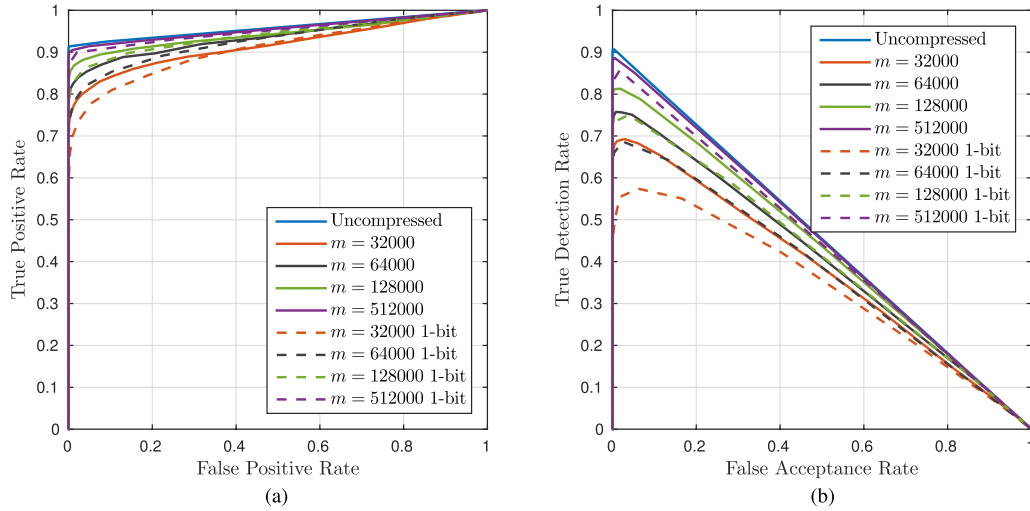


Fig. 3. ROC curves for the Dresden database. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

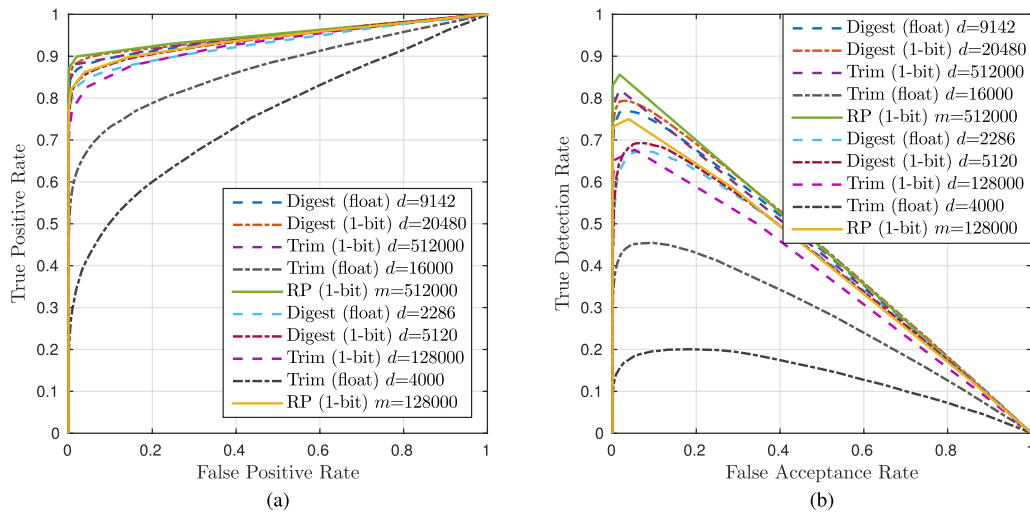


Fig. 4. ROC curves for the Dresden database. Binary random projections versus floating point and binary digest. Budget: 512000 bits, $m = 512000$, $d = 9142$ float, $d = 20480$ binary. Budget: 128000 bits, $m = 128000$, $d = 2285$ float, $d = 5120$ binary. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

100 test fingerprints per camera (mean correlation coefficient $\rho = 0.02$), generated as vectors of i.i.d. entries with standard Normal distribution. This simulated database confirms

the results obtained for the Dresden database, with binary random projections outperforming the digest technique. A final remark should be made about the matching complexity of the




Fig. 5. ROC curves for a synthetic Gaussian database. Binary random projections versus floating point and binary digest. Budget: 192000 bits, $m = 192000$, $d = 3429$ float, $d = 7680$ binary. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

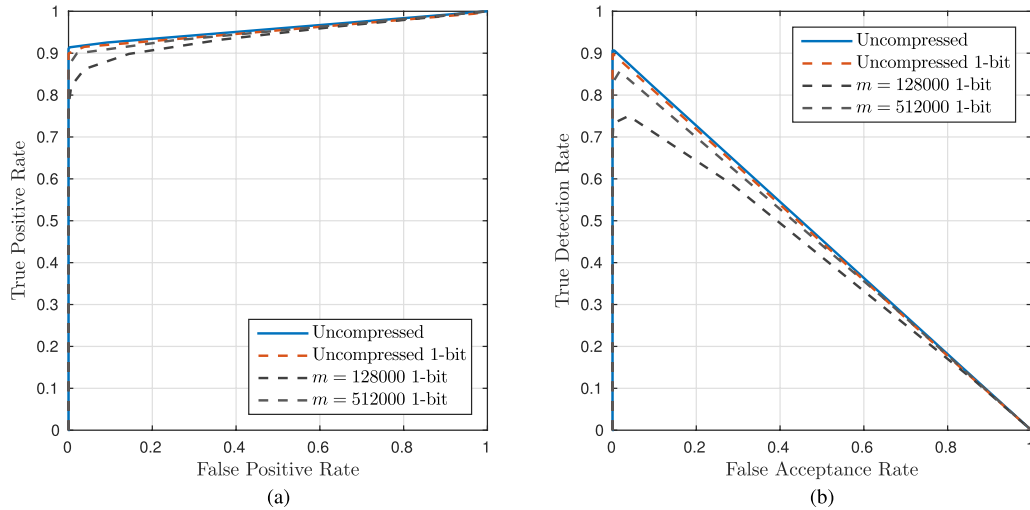


Fig. 6. ROC curves for the Dresden database. Binarization of the fingerprints is compared against binarized random projections. An 18-fold reduction of storage space is achieved with negligible performance loss by the $m = 512000$ case. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

digest technique. It would appear that since the digest has fewer coefficients, it could provide faster matching even though it is not as efficient as random projections in terms of compression ratio. This is not entirely true because it does not consider issues related to location information. A query fingerprint must be subsampled at the locations stored in the database, but such locations are different for each fingerprint. This means that the correctly subsampled version of the query must be assembled for every entry in the database and this either implies serialization of the matching process or a multiplication of the memory requirements equal to the degree of parallelism. Moreover, in the case of the binary digest, assembling the subsampled binary query has an even subtler problem of accessing non-contiguous bits while the smallest addressable unit of memory is typically a byte, thus causing additional overhead.

D. 1-bit Compression

Figures 2 and 3 also report the performance with binary quantization of random projections. In Section IV we have

given the theoretical reasons behind the good performance of 1-bit measurements. It is experimentally verified that binary random projections have good performance. The gap with respect to real-valued measurements is small compared to the significant savings in terms of storage and complexity of the matching operation. We experimentally observed that a system with m binary random projections typically shows a ROC nearly overlapped to the ROC of a system with $m/2$ real-valued measurements. Hence, as a rule of thumb we can consider a factor of 2 penalty in the dimension of the measurement space when using binary projections. However, storage requirements reduce by a factor of 64 (in case of double-precision measurements), so binary random projections exhibit extremely competitive performance. Bayram *et al.* proposed to compress fingerprints by binarization of the values [12]. This can be regarded as a limit case of the method proposed in this paper when the sensing matrix is the identity. Figure 6 shows that binary random projections yield further compression at a small expense in terms of performance. In particular, the test shown in the figures shows that the




Fig. 7. ROC curves for the Dresden database. Uncompressed detectors described in Section III-B are compared. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

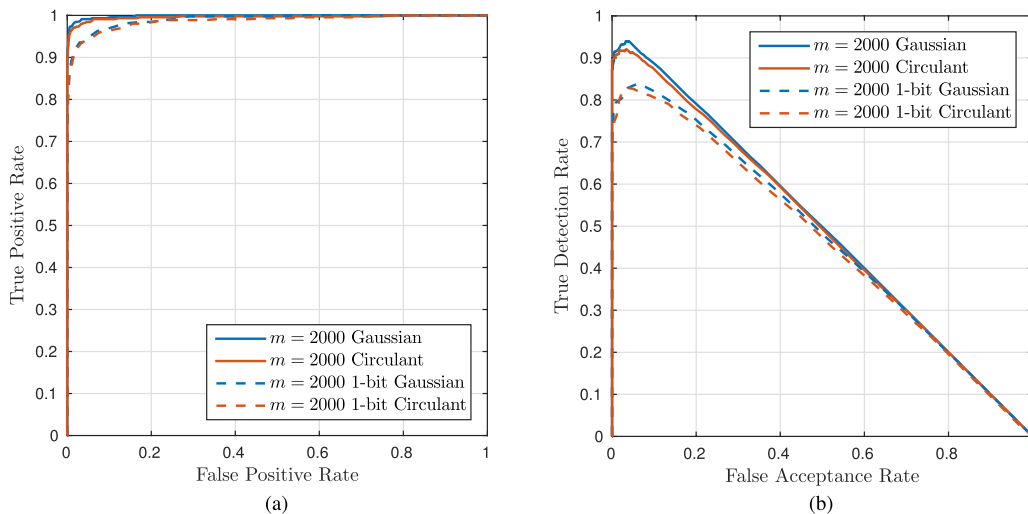


Fig. 8. PoliTo database ROC. Comparison between sensing matrices with i.i.d. Gaussian entries and circulant matrices and i.i.d. Gaussian entries in the first row. (a) True Positive Rate vs. False Positive Rate. (b) True Detection Rate vs. False Acceptance Rate.

decrease in the performance is negligible while the storage requirements is reduced by 18 times with respect to binary uncompressed fingerprints.

E. Suboptimal Detectors

We compare different types of detectors for the camera identification problem. As explained in Section III-B the optimal detector correlates the uncompressed noise residual of the test image with the uncompressed fingerprint modulated by the test image. As this cannot be mapped to the compressed domain, we investigate the suboptimality of the uncompressed detectors (17) and (18). From Fig. 7, we can notice that in practice all the detectors exhibit the same performance thus we settled on using detector (17), being the least computationally expensive.

F. Circulant vs. Gaussian Sensing Matrices

Several results, both theoretical and experimental [15], [33], [34], suggest that partial circulant matrices are

almost as effective as fully random Gaussian matrices, despite their structure and limited randomness. We compare the ROC obtained on the PoliTO database for Gaussian and circulant matrices, having the first row drawn as Gaussian i.i.d.. Experimental results shown in Fig. 8 confirm that circulant constructions perform very closely to the fully random ones, though they provide enormous advantages in terms of memory and computational requirements.

VI. CONCLUSIONS

This paper proposed a technique to address the issues of storage and matching complexity in camera fingerprint databases, by using random projections. Motivated by the incoherent nature of fingerprints based on PRNU patterns of camera sensors, we showed that random projections can effectively preserve the geometry of the database and significantly reduce the dimension of the problem with small penalties. We characterized the usage of real-valued and binary random measurements from a theoretical point of view in terms of the detection and false alarm probabilities.

Experimental tests have confirmed the validity of the proposed method on two databases of actual photographs. Practical issues such as the complexity of calculating random projections are of significant importance when dealing with million-pixel images, but we solved them by using circulant sensing matrices. The use of random projections for compression of camera fingerprints paves the way to many interesting applications involving maintaining large databases of fingerprints or applications requiring transmission of fingerprints over bandlimited channels. From this perspective, random projections are significantly better than the other existing methods discussed in this paper because they can provide higher compression ratios and improved scalability, *i.e.*, a fine-grained control over the compression/performance tradeoff by modulating the number of projections according to the specific needs, and an embedded representation where a compressed version of the fingerprint already embeds versions at higher compression ratios (fewer measurements used).

REFERENCES

- [1] J. Fridrich, "Digital image forensics," *IEEE Signal Process. Mag.*, vol. 26, no. 2, pp. 26–37, Mar. 2009.
- [2] J. Lukáš, J. Fridrich, and M. Goljan, "Determining digital image origin using sensor imperfections," *Proc. SPIE*, vol. 5685, pp. 249–260, Apr. 2005.
- [3] J. Lukáš, J. Fridrich, and M. Goljan, "Digital camera identification from sensor pattern noise," *IEEE Trans. Inf. Forensics Security*, vol. 1, no. 2, pp. 205–214, Jun. 2006.
- [4] M. Chen, J. Fridrich, M. Goljan, and J. Lukáš, "Determining image origin and integrity using sensor noise," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 74–90, Mar. 2008.
- [5] M. Goljan and J. Fridrich, "Camera identification from cropped and scaled images," *Proc. SPIE*, vol. 6819, pp. 68190E-1–68190E-13, 2008.
- [6] S. Bayram, H. T. Sencar, and N. Memon, "Efficient techniques for sensor fingerprint matching in large image and video databases," *Proc. SPIE*, vol. 7541, pp. 754109-1–754109-8, Jan. 2010.
- [7] M. Goljan, J. Fridrich, and T. Filler, "Managing a large database of camera fingerprints," *Proc. SPIE*, vol. 7541, pp. 754108-1–754108-12, Jan. 2010.
- [8] Y. Hu, B. Yu, and C. Jian, "Fast camera fingerprint search algorithm for source camera identification," in *Proc. 2nd Int. Conf. Comput. Sci. Appl. (CSA)*, Dec. 2009, pp. 1–5.
- [9] Y. Hu, C.-T. Li, Z. Lai, and S. Zhang, "Fast camera fingerprint search algorithm for source camera identification," in *Proc. 5th Int. Symp. Commun. Control Signal Process. (ISCCSP)*, May 2012, pp. 1–5.
- [10] Y. Hu, C.-T. Li, and Z. Lai, "Fast source camera identification using matching signs between query and reference fingerprints," in *Multimedia Tools and Applications*. New York, NY, USA: Springer-Verlag, 2014, pp. 1–24. [Online]. Available: <http://dx.doi.org/10.1007/s11042-014-1985-3>
- [11] M. Goljan and J. Fridrich, "Sensor fingerprint digests for fast camera identification from geometrically distorted images," *Proc. SPIE*, vol. 8665, pp. 86650B-1–86650B-10, Mar. 2013.
- [12] S. Bayram, H. Sencar, and N. Memon, "Efficient sensor fingerprint matching through fingerprint binarization," *IEEE Trans. Inf. Forensics Security*, vol. 7, no. 4, pp. 1404–1413, Aug. 2012.
- [13] W. B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space," *Contemp. Math.*, vol. 26, pp. 189–206, 1984.
- [14] D. Achlioptas, "Database-friendly random projections: Johnson–Lindenstrauss with binary coins," *J. Comput. Syst. Sci.*, vol. 66, no. 4, pp. 671–687, 2003.
- [15] J. Vybíral, "A variant of the Johnson–Lindenstrauss lemma for circulant matrices," *J. Funct. Anal.*, vol. 260, no. 4, pp. 1096–1105, 2011.
- [16] L. Jacques, J. N. Laska, P. T. Boufounos, and R. G. Baraniuk, "Robust 1-bit compressive sensing via binary stable embeddings of sparse vectors," *IEEE Trans. Inf. Theory*, vol. 59, no. 4, pp. 2082–2102, Apr. 2013.
- [17] M. Goljan, M. Chen, and J. Fridrich, "Identifying common source digital camera from image pairs," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, vol. 6, Sep. 2007, pp. VI-125–VI-128.
- [18] D. L. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [19] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [20] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Construct. Approx.*, vol. 28, no. 3, pp. 253–263, 2008.
- [21] P. Indyk and R. Motwani, "Approximate nearest neighbors: Towards removing the curse of dimensionality," in *Proc. 13th Annu. ACM Symp. Theory Comput.*, 1998, pp. 604–613.
- [22] A. Andoni and P. Indyk, "Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions," in *Proc. 47th Annu. IEEE Symp. Found. Comput. Sci. (FOCS)*, Oct. 2006, pp. 459–468.
- [23] M. Datar, N. Immorlica, P. Indyk, and V. S. Mirrokni, "Locality-sensitive hashing scheme based on p -stable distributions," in *Proc. 20th Annu. Symp. Comput. Geometry*, 2004, pp. 253–262.
- [24] A. Hinrichs and J. Vybíral, "Johnson–Lindenstrauss lemma for circulant matrices," *Random Struct. Algorithms*, vol. 39, no. 3, pp. 391–398, Oct. 2011.
- [25] H. Rauhut, "Circulant and Toeplitz matrices in compressed sensing," in *Proc. Signal Process. Adapt. Sparse Struct. Represent. (SPARS)*, 2009, pp. 1–6.
- [26] J. Fridrich and M. Goljan. (Jan. 2010). "Derivation of ROCs for composite fingerprints and sequential trimming." Dept. Elect. Comput. Eng., Binghamton Univ., Binghamton, NY, USA, Tech. Rep. [Online]. Available: <http://www.ws.binghamton.edu/fridrich/Research/rocs.pdf>
- [27] S. S. Vempala, *The Random Projection Method*, vol. 65. Providence, RI, USA: AMS, 2004.
- [28] F. Krahermer and R. Ward, "New and improved Johnson–Lindenstrauss embeddings via the restricted isometry property," *SIAM J. Math. Anal.*, vol. 43, no. 3, pp. 1269–1281, 2011.
- [29] T. Gloe and R. Böhme, "The Dresden image database for benchmarking digital image forensics," *J. Digit. Forensic Pract.*, vol. 3, nos. 2–4, pp. 150–159, 2010. [Online]. Available: <http://forensics.inf.tu-dresden.de/ddimgdb/>
- [30] M. Goljan, J. Fridrich, and T. Filler, "Large scale test of sensor fingerprint camera identification," *Proc. SPIE*, vol. 7254, pp. 72540I-1–72540I-12, Feb. 2009.
- [31] University of Binghamton Digital Data Embedding Laboratory. *Camera Fingerprint Software*. [Online]. Available: http://dde.binghamton.edu/download/camera_fingerprint/, accessed Jan. 2014.
- [32] T. Gloe, S. Pfennig, and M. Kirchner, "Unexpected artefacts in PRNU-based camera identification: A 'Dresden image database' case-study," in *Proc. Multimedia Secur.*, 2012, pp. 109–114.
- [33] J. Haupt, W. U. Bajwa, G. Raz, and R. Nowak, "Toeplitz compressed sensing matrices with applications to sparse channel estimation," *IEEE Trans. Inf. Theory*, vol. 56, no. 11, pp. 5862–5875, Nov. 2010.
- [34] W. Yin, S. Morgan, J. Yang, and Y. Zhang, "Practical compressive sensing with Toeplitz and circulant matrices," *Proc. SPIE*, vol. 7744, pp. 77440K-1–77440K-10, Jul. 2010.



Diego Valsesia received the M.Sc. degree in telecommunications engineering from the Politecnico di Torino, Turin, Italy, and the M.Sc. degree in electrical and computer engineering from the University of Illinois at Chicago, Chicago, IL, in 2012. He is currently pursuing the Ph.D. degree with the Department of Electronics and Telecommunications, Politecnico di Torino. His main research interests include compression of remote sensing images, compressed sensing, and sparse representations.



Giulio Coluccia (M'12) received the B.Sc. degree and the M.Sc. degree in telecommunications engineering from the Politecnico di Torino, Turin, Italy, in 2003 and 2005, respectively, and the Ph.D. degree in electronic and communications engineering from the Electronics Department, Politecnico di Torino, in 2009, under the supervision of Prof. G. Taricco.

He is currently a Post-Doctoral Researcher with the Image Processing Laboratory, Politecnico di Torino, led by Prof. E. Magli. His research is focused on compressed sensing, with a particular interest in its application to image processing and forensics, multidimensional signals, and to distributed source coding and wireless sensor networks. He is involved in a 5-year ERC project entitled CRISP-Toward Compressive Information Processing Systems funded by the European Research Council.



Enrico Magli (S'97–M'01–SM'07) received the M.Sc. and Ph.D. degrees from the Politecnico di Torino, Turin, Italy, in 1997 and 2001, respectively. He is currently an Associate Professor with the Politecnico di Torino. His research interests are in the field of compressive sensing, image and video coding, and vision. He was a corecipient of the IEEE Geoscience and Remote Sensing Society Transactions Prize Paper Award in 2011. He is an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY, the IEEE TRANSACTIONS ON MULTIMEDIA, and the *EURASIP Journal on Image and Video Processing*, and an IEEE Distinguished Lecturer from 2015 to 2016.



Tiziano Bianchi (S'03–M'05) received the M.Sc. degree (Laurea) in electronics engineering and the Ph.D. degree in information and telecommunication engineering from the University of Florence, Italy, in 2001 and 2005, respectively.

He was with the Department of Electronics and Telecommunications, University of Florence, from 2005 to 2012, as a Research Assistant. Since 2012, he has been with the Department of Electronics and Telecommunications, Politecnico di Torino, as an

Assistant Professor. His research interests have involved signal processing in communications and processing of Synthetic Aperture Radar images. He has authored over 100 papers in international journals and conference proceedings. His current research topics include multimedia security technologies, signal processing in the encrypted domain, and security aspects of compressed sensing.