The identification of mobile phones through the fingerprints of their built-in magnetometer: an analysis of the portability of the fingerprints

Gianmarco Baldini
Directorate for Space, Security and Migration
European Commission, Joint Research Centre (JRC)
Ispra, Italy
gianmarco.baldini@ec.europa.eu

Irene Amerini
Media Integration and Communication Center (MICC)
University of Florence
Florence, Italy
irene.amerini@unifi.it

Gary Steri
Directorate for Space, Security and Migration
European Commission, Joint Research Centre (JRC)
Ispra, Italy
gary.steri@ec.europa.eu

Roberto Caldelli
Media Integration and Communication Center (MICC)
University of Florence
Florence, Italy
roberto.caldelli@unifi.it

***Abstract*—The identification of mobile phones through their built-in electronic sensors has been proven in literature for cameras, microphones and accelerometers demonstrating to have various applications in security, forensics or in the fight against the counterfeiting of electronic products. The identification of a sensor (and consequently of the mobile phone) is possible through the exploitation of small but significant differences in the physical components of the sensor itself. These physical differences are mainly generated during the production process of the sensor in the manufacturing plant producing small but reproducible variations in the digital output generated by the built-in sensor. In particular, in this paper, we investigate the identification of mobile phones through the built-in magnetometer sensor, which has received very limited attention by the research community so far. In particular, the specific aspects of robustness and portability of such a fingerprint have been analyzed. Different stimulation conditions, diverse kinds of features and classification procedures have been considered achieving very promising results.**

I. Introduction

The identification of mobile phones through their built-in components can be performed by exploiting the tiny but significant differences existing in the composing materials or introduced during the manufacturing process. These differences result in small and unperceived distortions in the digital output generated by the mobile phone. They are also called *fingerprints* of the component, similarly to the fingerprints of a human being. For example, the digital camera of a mobile phone introduces in every image processed and stored in the memory of the phone, a specific feature which is called Sensor Pattern Noise (SPN). The SPN can be extracted from a sufficient number of images as demonstrated in [1]. In a similar way, researchers have demonstrated the capability to identify mobile phones with different degrees of accuracy from the built-in accelerometers, radio frequency components, microphones and so on. The identification of mobile phones through their components has various applications in security, forensics or fight against the counterfeiting of electronic products. In fact, identification proofs based on physical characteristics are much more difficult to be faked and reproduced because they are intrinsically related to the component and the mobile phone itself. Such fingerprints can be used to perform multi-factor authentication where physical identification is combined with cryptographic authentication (see [2]). In this kind of application, it is important that a unique response is systematically generated by a specific mobile phone component often in correspondence to a specific stimulus. In particular, a two-fold purpose has to be obtained: to identify phones of the same model but different serial numbers (*intra-model* identification) and phones of different models and brands (*inter-model* identification). The former is usually more difficult to be achieved compared to the latter, essentially because phones of the same model are built using the same components, while in different models we can find different materials and components.

In this paper, we describe an approach for magnetometers fingerprinting based on the stimulation of such a sensor through magnetic fields generated by a solenoid connected to a generic sound board of a computer. The proposed method involves the extraction of statistical features, both in the temporal and in the spectral domain, for each phone to generate a set of fingerprints; such features are calculated onto the output data acquired by the magnetometer. Furthermore, we focus on the specific aspect of robustness and portability of the fingerprints by analyzing different stimulating sound boards, devising two

This paper has been partially supported by the project ESPRESS (Smartphone identification based on on-board sensors for security applications) co-funded by Fondazione Cassa di Risparmio di Firenze (Italy) within the Scientific Research and Technological Innovation framework.

978-1-5386-1585-0/17/$31.00

different sets of features to build the fingerprint and evaluating various classification methods.

The structure of this paper is the following: Section II provides an overview of the literature of electronic device fingerprinting in mobile phones, Section III describes the methodology introducing, in particular, the approach used to stimulate the magnetometer, the extraction of the statistical features from the digital output and the application of the machine learning algorithms to perform the classification process. Then, Section IV presents the experimental results on a set of nine phones including phones of the same model to also test the *intra-model* identification capability, and finally Section V concludes the paper.

II. Related Works

The fingerprint concept has been applied to different components of the mobile phones. A recent survey [3] has highlighted many components that can be used to identify mobile phones. One of the first example is based on the possibility to detect the digital camera of a mobile phone thanks to the imperfections present in various components like: lenses, Color Filter Array (CFA), the CCD/CMOS sensor or even the software to process the images before storing them (e.g., compression algorithms). The SPN presented in [1] is the most adopted and referenced approach because of its high accuracy also in the *intra-model* case and for its persistence in time. The SPN is based on the non-uniformity of each sensor pixel sensitivity to light which determines that a sistematic noise pattern is embedded within each digital content the sensor acquires [4].

Fingerprinting of the radio frequency components of wireless devices has also been performed for different wireless standards. Researchers have applied Radio Frequency (RF) fingerprinting to WLAN in [5], to Global System for Mobile Communications (GSM) in [6] and to ZigBee devices [7]. In all cases, the fingerprinting technique is based on the selection of statistical features of the collected and processed radio frequency signals emitted the mobile phones or other wireless devices. This process is also called Specific Emitter Identification (SEI) [8], [9].

The mobile phone's microphone is another component that can be used for classification. In [10], large-size raw feature vectors are obtained by averaging the log-spectrogram of a speech recorded along the time axis for each mobile phone. The authors focused on inter-model identification; the features were fed to three distinct classifiers: the Sparse representations Classifier (SRC), the Support Vector Machine (SVM) and Nearest Neighbour (NN). SVM provided the best performance in many configurations.

Finally, accelerometers and gyroscopes can also be used for fingerprinting as demonstrated in [11], [12] and [13]. The mobile phones are submitted to repeatable motion patterns, which stimulate the accelerometer and gyroscope to produce specific responses. The small variations in the responses are used to create the fingerprints and distinguish among mobile phones. Both *inter-model* and *intra-model* identification were performed with very good accuracy. In [14], gyroscope and accelerometer sensors are used in combination with features extracted from the SPN to achieve a robust fingerprint to classify different smartphones. A very recent paper [15] extends this work investigating the possibility to identify mobile phones through the combination of features coming from different sensors: accelerometer, gyroscope, magnetometer and microphone-speaker. In this case the authors have not stimulated the magnetometer sensor but have only relied on the natural magnetic field of Earth. Even if this can be considered a valid method in many applications, some specific others (e.g., fight against counterfeiting) require that the fingerprint is correlated to a specific stimulus, which can be generated artificially (e.g., as in challenge-response authentication) and it is not dependent on natural phenomena (e.g., because this would limit the range of potential stimuli). In another recent paper [16], the authors have generated an external magnetic field stimulus by positioning a magnet in front of a low cost rotating platform, where the mobile phone is mounted in a precise position to stimulate the magnetometer sensor. Both in [15], [16] it is demonstrated that magnetometer sensor can be used to exploit an effective fingerprint to distinguish different smartphones. In this paper we explore in depth this fact analysing the magnetometer behavior under different stimulations and using a diverse implementation of the fingerprint construction.

III. Methodology

The setup of the test bed used to generate the magnetic field stimuli and to collect the response through a smartphone is shown in Figure 1.

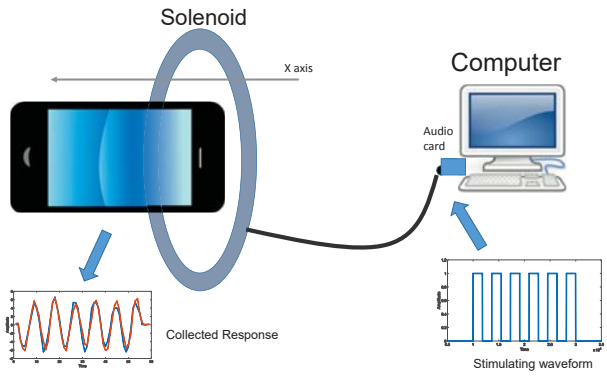

Fig. 1: Test setup to stimulate the magnetometer of a mobile phone.

In order to show the feasibility of mobile phone identification using low cost equipment, the magnetometer is stimulated using a cost-effective solenoid, which is connected to a consumer mass market audio board. The audio board runs a repeated set of 260 synthetic waveforms based on square waves. The waveform is shown in Figure 2 and it is defined according to the following considerations: a) a sharp impulse is needed to stimulate the magnetometer, b) the distance between the square waves is defined on the average hysteresis values of the common mass market magnetometers and then

empirically tested. Others waveforms might be more suitable to this purpose and this aspect will be investigated as future work. The responses of mobile phones in correspondence to a waveform stimulation are shown in Figure 3.

In particular three different sound cards have been used to generate the magnetic field stimulus. This is one of the objectives of this paper: to investigate if different sound cards produce a different fingerprint thus enabling the identification process and what characteristics of the sound cards can be exploited to improve the identification accuracy.

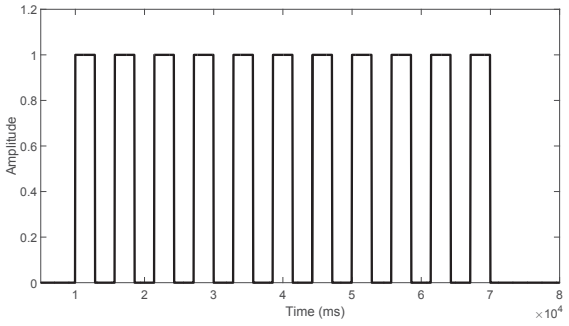

Fig. 2: Waveform used to stimulate the magnetometer of the mobile phone.

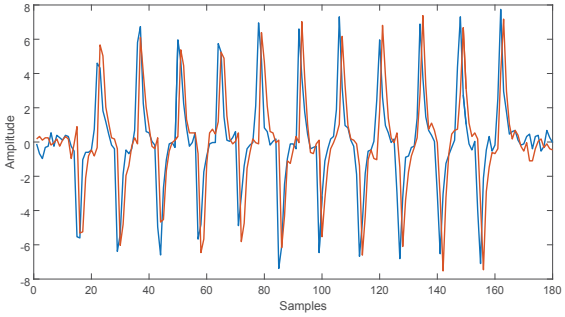

Fig. 3: Different responses of the built-in magnetometer for two smartphones to the stimula in Figure 2.

A. Workflow of the proposed method

In the following, the main steps of the proposed method to identify smartphones on the basis of magnetometer fingerprints are detailed.

Each mobile phone (in the set of N phones) has been submitted to the magnetic field generated by the solenoid for all the three different sound cards, as described before. The digital output response of the magnetometer has been collected using a free available mobile application called *AndroSensor*[1] (the same version of the app is installed in all phones). In this paper only phones supporting the Android operating system are considered but a similar study can be obviously performed on iOS phones. The collection of the data from the mobile phones was set to a frequency of 20 Hz (i.e., samples were taken by the magnetometer every 50 ms). This value was chosen because it was the lower common limit for all the mobile phones sensors in the dataset. The responses from each phone and for each sound card are synchronized and normalized as in [2], [11] and [6].

Then statistical features are extracted from the responses collected from each phone. A complete description of the employed features are detailed in the following Section III-B. In particular, two set of features have been considered in the proposed analysis named *Features set1* and *Features set2*. Moreover a set of three different supervised classifiers have been applied to the extracted statistical features in order to achieve device identification: the Support Vector Machine (SVM) and the K-Nearest Neighbour (K-NN) classifiers are trained using the *Features set1*, while the Bagged Decision Tree is used with the *Features set2* (for a detailed description of classifiers see Section III-C).

Once the system has been properly trained, each time a smartphone needs to be identified in the test phase, a sensor reading is collected; the *Features set1* and *Features set2* are computed respectively for the two considered situations and then the identification phase takes place.

B. Features extraction

1) Features set1: The first set of considered features is listed in Table I; they are 6 measures computed in the time domain and in the frequency domain both for amplitude and phase of the Fast Fourier Transform (FFT). Then, *Features set1* includes a total of 18 features (6x3). The raw values, acquired by the magnetometer sensor, are collected by the *AndroSensor* application only along the x axis of the smartphone (the main axis where the solenoid and the smartphone magnetometer are aligned as described in Figure 1) at a certain time stamp t: $m(t) = m_x(t)$.

Features set1
Shannon Entropy
Log Energy Entropy
Std-Deviation
Variance
Skewness
Kurtosis

TABLE I: Time and frequency domain features used for *Feature set1*.

2) Features set2: The raw values, acquired by the magnetometer sensor, are collected again by the *AndroSensor* application, but in this case, along all the three axes of the smartphone at a certain time-stamp t: $m(t) = (m_x(t), m_y(y), m_z(t))$. The features are extracted in both time and frequency domains by using the MIRToolbox [17] starting from the following signal:

$$|m(t)| = \sqrt{m_x^2(t) + m_y^2(t) + m_z^2(t)}$$

[1]Downloadable from Google Play for Android.

In total a vector of 21 features (see Table II), consisting of 10 temporal and 11 spectral features (only amplitude is taken), compose the *Features set2*.

Features set2	
Time Features	*Frequency Features*
Mean	Spectral Spread
Std-Deviation	Spectral Centroid
Average Deviation	Spectral Skewness
Skewness	Spectral Kurtosis
Kurtosis	Entropy
RMS	Flatness
Max	Roll Off
Min	Roughness
ZCR	Irregularity
Non-negative count	Spectral RMS
	Low Energy Rate

TABLE II: Time and frequency domain features used for *Features set2*.

C. Classification methods

1) SVM and K-NN Classifiers: The SVM is a very well known technique in supervised machine learning and it has been used in this paper since it has demonstrated its effectiveness for RF fingerprinting in the literature (see [18], [10] and [13]). Two parameters must be optimized in the application of SVM: the scaling factor and the box constraint [19]. A grid approach based on the metric of the overall classification accuracy is used to calculate the optimum values. In the experimental tests, the scaling factor equals to $2^7(128)$ and the box constraint is set at 2^{22} (4194304). The One Against One (OAO) approach has been used for multi-classification based on SVM. Furthermore, the K-NN technique has also been considered with different values of the K neighbor distance. A 10-fold partition has been adopted for training and classification both for SVM and K-NN techniques. In the 10-fold method, each collection of 260 responses (one for each input stimulus) is divided into ten blocks. Nine blocks are used for training (234 responses) and one block is held out for classification (26 responses). The training and testing phases are repeated ten times until each of the ten blocks has been held out and classified. The final cross-validation performance statistics are calculated by averaging the results of all folds. In this way, the presence of a bias in a specific training set is averaged or mitigated.

2) Bagged Decision Tree Classifier: The other supervised classifier taken into consideration is the Random Forest classifier similarly as in [15]. The Random Forest classifier is based on a Decision Tree model built following the ID3 algorithm (Iterative Dichotomiser 3), and its structure represents the set of decisions that allows to classify an input feature vector. Each branch of the structure represents an interval of the value for a selected feature. The set of the taken decisions (each for every feature) guides the execution through a path in the tree structure that leads to a *leaf node*, which finally indicates a label for the input feature vector. The idea behind this classification technique is to build a structure that, given an input test feature vector, allows to output a label based on the value of each feature evaluating its affinity with a certain value interval. Techniques like bagging and boosting have been introduced to improve the classification accuracy of a Decision Tree model limiting the variance of different *weak* learners and reducing the correlation among the trees that compose the forest (see [20] for details).

IV. Experimental results

In this section some of the different experimental tests that have been carried out are presented to demonstrate the reliability in smartphone classification of the technique illustrated above. A set of nine mobile phones (some of them of the same brand and model) are stimulated using three consumer mass market audio boards. Each phone is stimulated with a repetition of 260 waveforms and the magnetometer response is collected and processed in order to extract the two sets of the previous described features. The goal of this experimental campaign is to understand if the identification accuracy obtained with one sound card is comparable and coherent with the accuracy obtained with another one. On the other side, it is also interesting to evaluate if a sound card of better quality is able to provide an improved identification accuracy compared to a lower quality one. Furthermore, we want to investigate if the classification accuracy is independent from the generation of the fingerprint and from the classifiers that was used. The three employed sound cards belong to two desktop PCs and one laptop computer. In detail, the first sound card (RealTek high definition audio), called SC1, belongs to a laptop of type Fujitsu Celsius, the SC2 is a PC desktop computer sound card (model IDT 92HD94) different from SC1 but of similar quality and specifications. The third one (SC3) is another PC desktop sound card (model Xonar Essence STX) of better quality and performance than the previous ones.
In order to evaluate the portability of the illustrated method, various tests have been defined on the two different features sets and on the basis of the considered classifiers. In Table III the 9 smartphones composing the dataset are listed. To better simulate real operative conditions, some of the selected devices are identical (i.e., with same brand and model: three HTC One, two Samsung Galaxy S5 and two Sony Experia). The obtained results have been evaluated in terms of *F-score* ($F1$), *Precision* and *Recall* which are defined as in Equation (1):

$$F1 = 2 * \left(\frac{Pr * Re}{Pr + Re} \right) = \frac{2 * TP}{2 * TP + FN + FP} \qquad (1)$$

where $Pr = \frac{TP}{(TP+FP)}$ and $Re = \frac{TP}{(TP+FN)}$ stands for *Precision* and *Recall* respectively.

A. Results on Features set1 with SVM and K-NN

Performances regarding the first configuration of features (*Features set1*) and the corresponding two classifiers are presented in Table IV, which shows the results in terms of *F-score* ($F1$) among the three different sets of data obtained from

Smartphone model	Quantity	Index
HTC One	3	1-3
Huawei Ascend Mate	1	4
Samsung Galaxy S5 Android version 4.4	1	5
Samsung Galaxy S5 Andriod version 6.0	2	6-7
Sony Experia	2	8-9

TABLE III: List of mobile phones used in the experiments.

SVM	*F-score* ($F1$)
SC1	0.741
SC2	0.685
SC3	0.877
KNN (K=5)	***F-score*** ($F1$)
SC1	0.617
SC2	0.633
SC3	0.808
KNN (K=10)	***F-score*** ($F1$)
SC1	0.613
SC2	0.663
SC3	0.901
KNN (K=15)	***F-score*** ($F1$)
SC1	0.614
SC2	0.656
SC3	0.813

TABLE IV: *F-score* ($F1$) obtained with different sound cards; *Features set1*, SVM and K-NN classifiers.

the three sound cards. Table IV shows the overall performance using the SVM classifier and K-NN classifier of order 5, 10 and 15. From the results, it can be seen that the identification accuracy is similar for sound cards of similar quality (SC1 and SC2), while a better sound card (SC3) produces a superior result. The results for K-NN, though generally worst than SVM (anyway this should have been weighted with a larger computational time), witness, even much more, the improvement of the accuracy for the SC3 case. To summarize, the use of different sound cards seems to have an impact on the identification and this fact suggests that a sound card of better quality could be a tuning factor for smartphone classification purpose.

B. Results on Features set2 with Bagged Decision Tree

In this section, the experimental conditions were changed in order to highlight the actual distinctiveness potentialities of the magnetometer beyond the approach used both for feature extraction phase and for the successive classification.
The acquired magnetometer signal (i.e., the total of the 260 responses collected from each of the 9 phones in the dataset of Table III) is subdivided and separately used for training and testing phases. The response to each waveform is gathered together and appended one after another creating a unique signal. Different numbers of samples have been proven to extract the fingerprint, and at the end, 90 samples for each smartphone reading are used to train the classifier. Each sample is obtained by processing a non-overlapped chunk of 10 seconds of the signal (i.e., 900 seconds are used for training). The remaining part of the signal is used as test set. The first proposed experiment takes in exam the classification performance respect to the three considered audio boards with

	F1	Precision	Recall
SC1	0.891	0.926	0.857
SC2	0.868	0.879	0.858
SC3	0.907	0.950	0.867

TABLE V: Results in terms of *F-score* ($F1$), *Precision* and *Recall* on the three considered audio boards: *Features set2*, Bagged Decision Tree classifier.

the Bagged Decision Tree classifier trained on the *Features set2* . The obtained results are shown in Table V; it can be seen that the performances are good for all the three different kind of audio boards with a *F-score* ($F1$) of about 88% on average demonstrating the robustness of the devised fingerprint.

Furthermore, in Figure 4 the confusion matrix of one of the previous case (SC2 audio board) is reported for an in-depth analysis. It is interesting to note that the classifier fails to predict the right class when it encounters a smartphone of the same brand and model (see Figure 4 where classes 1-2 correspond to HTC One and 8-9 to Sony Experia), i.e., during intra-model identification. The prediction it is almost perfect in all the other cases.

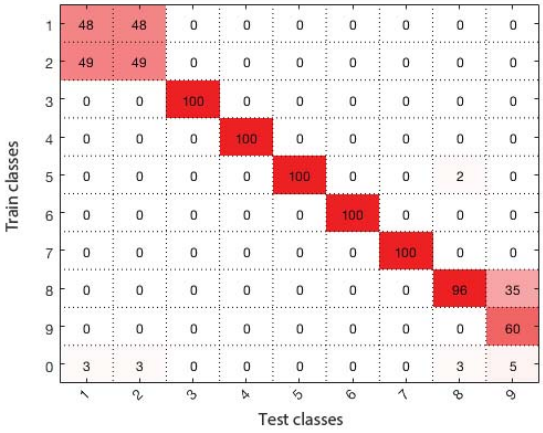

Fig. 4: Confusion matrix for the SC2 audio board in Table V (values expressed as %). '1-9' classes correspond to the smartphones in Table III, '0' is the alien class i.e., the test sensor reading is not recognized as belonging to one of the smartphone in the training set.

In the last experiment, we have evaluated the case in which the training set is performed on an audio board (i.e., SC1 laptop audio board) and the test set is composed by the reading of another sound card (i.e., SC2) and viceversa. In Table VI the results in term of *F-score* ($F1$) and *Precision/Recall* are reported demonstrating a very good portability behavior. The results related to the others combination of sound cards are straightforward.

V. Conclusions

This paper has proposed a methodology based on different classifiers to evaluate the performance and portability of the fingerprints obtained from the magnetometer sensor of a

Train	Test	F1	Precision	Recall
SC1	SC2	0.866	0.884	0.848
SC2	SC1	0.868	0.887	0.849

TABLE VI: Results on portability in terms of *F-score (F1)*, *Precision* and *Recall*; *Features set2*, Bagged Decision Tree classifier.

mobile phone exploiting two different sets of features and three different sounds cards used to generate the magnetic stimuli through a solenoid. The presented technique has been tested on a dataset composed by 9 phones of different brand and models and the obtained results demonstrated a very good portability behavior of such fingerprints and the reliability of the proposed procedure. In particular, the *Features set2* in combination with the Bagged Decision Tree classifier and the use of all the three magnetic axes has demonstrated to perform better respect to the other configuration (*Features set1*, SVM or K-NN) where only one magnetic axis (i.e., x axis) is used, suggesting to be more suitable for the classification task. Future works will be devoted to increase the number of the considered smartphones in the dataset and also evaluating different waveforms to stimulate the magnetometers.

REFERENCES

[1] J. Lukas, J. Fridrich, and M. Goljan, "Digital camera identification from sensor pattern noise," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 205–214, June 2006.

[2] W. C. Suski II, M. A. Temple, M. J. Mendenhall, and R. F. Mills, "Using spectral fingerprints to improve wireless network security," in *IEEE GLOBECOM 2008-2008 IEEE Global Telecommunications Conference*. IEEE, 2008, pp. 1–5.

[3] G. Baldini and G. Steri, "A survey of techniques for the identification of mobile phones using the physical fingerprints of the built-in components," *IEEE Communications Surveys Tutorials*, vol. PP, no. 99, pp. 1–1, 2017.

[4] J. Fridrich, "Digital image forensics," *IEEE Signal Processing Magazine*, vol. 26, no. 2, pp. 26–37, March 2009.

[5] G. Huang, Y. Yuan, X. Wang, and Z. Huang, "Specific emitter identification for communications transmitter using multi-measurements," *Wireless Personal Communications*, pp. 1–20, 2016. [Online]. Available: http://dx.doi.org/10.1007/s11277-016-3696-8

[6] D. R. Reising, M. A. Temple, and M. J. Mendenhall, "Improved wireless security for gmsk-based devices using RF fingerprinting," *International Journal of Electronic Security and Digital Forensics*, vol. 3, no. 1, pp. 41–59, 2010.

[7] T. J. Bihl, K. W. Bauer, and M. A. Temple, "Feature selection for rf fingerprinting with multiple discriminant analysis and using zigbee device emissions," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 8, pp. 1862–1874, Aug 2016.

[8] H. C. A. van Tilborg and S. Jajodia, Eds., *Specific Emitter Identification (SEI)*. Boston, MA: Springer US, 2011, pp. 1243–1243.

[9] J. Zhang, F. Wang, O. A. Dobre, and Z. Zhong, "Specific emitter identification via Hilbert-Huang Transform in single-hop and relaying scenarios," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 6, pp. 1192–1205, June 2016.

[10] C. L. Kotropoulos, "Source phone identification using sketches of features," *IET Biometrics*, vol. 3, no. 2, pp. 75–83, June 2014.

[11] S. Dey, N. Roy, W. Xu, R. R. Choudhury, and S. Nelakuditi, "AccelPrint: Imperfections of Accelerometers Make Smartphones Trackable," in *2014 Network and Distributed System Security (NDSS) Symposium*, Feb 2014.

[12] H. Bojinov, Y. Michalevsky, G. Nakibly, and D. Boneh, "Mobile device identification via sensor fingerprinting," *CoRR*, vol. abs/1408.1416, 2014. [Online]. Available: http://arxiv.org/abs/1408.1416

[13] G. Baldini, G. Steri, F. Dimc, R. Giuliani, and R. Kamnik, "Experimental identification of smartphones using fingerprints of built-in micro-electro mechanical systems (mems)," *Sensors*, vol. 16, no. 6, p. 818, 2016.

[14] I. Amerini, P. Bestagini, L. Bondi, R. Caldelli, M. Casini, and S. Tubaro, "Robust smartphone fingerprint by mixing device sensors features for mobile strong authentication," in *Media Watermarking, Security, and Forensics*. Ingenta, 2016, pp. 1–8.

[15] I. Amerini, R. Becarelli, R. Caldelli, A. Melani, and M. Niccolai, "Smartphone fingerprinting combining features of on-board sensors," *IEEE Transactions on Information Forensics and Security*, vol. PP, no. 99, pp. 1–1, 2017.

[16] G. Baldini, F. Dimc, R. Kamnik, G. Steri, R. Giuliani, and C. Gentile, "Identification of mobile phones using the built-in magnetometers stimulated by motion patterns," *Sensors*, vol. 17, no. 4, p. 783, 2017.

[17] O. Lartillot and P. Toiviainen, "MIR in Matlab: A toolbox for musical feature extraction from audio," in *International Society for Music Information Retrieval Conference (ISMIR)*, 2007,

[18] J. Hasse, T. Gloe, and M. Beck, "Forensic identification of GSM mobile phones," in *Proceedings of the first ACM workshop on Information hiding and multimedia security*. ACM, 2013, pp. 131–140.

[19] N. Cristianini and J. Shawe-Taylor, *An introduction to support vector machines and other kernel-based learning methods*. Cambridge university press, 2000.

[20] T. Ho, "The random subspace method for constructing decision forests," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 20, pp. 832–844, 1998.