

UNSUPERVISED FUSION FOR FORGERY LOCALIZATION EXPLOITING BACKGROUND INFORMATION

P. Ferrara¹, M. Fontani^{1,2}, T. Bianchi³, A. De Rosa^{1,2}, A. Piva^{1,2} and M. Barni^{2,4}

¹Dept. of Information Engineering, University of Florence, 50139, Firenze, Italy
Email: pasquale.ferrara@unifi.it

²National Inter-University Consortium for Telecommunications, Italy

³Dept. of Electronics and Telecommunications, Politecnico di Torino, Torino, Italy

⁴Dept. of Information Engineering and Mathematics, University of Siena, 53100, Siena, Italy

ABSTRACT

When image authenticity verification has to be carried out without any knowledge about the possible processing undergone by the image under analysis, it is fundamental to rely on a multi-clue approach, that merges the information stemming from several complementary forensic tools. This paper introduces a fully automatic framework for fusing the maps created by a set of unsupervised forgery localization algorithms, indicating possible manipulated areas. The framework takes into account the forgery maps, their reliability and the compatibility among the different traces considered by the tools. The achieved localization map is then refined by exploiting image content, thus improving the performance of the proposed system with respect to state of the art approaches.

Index Terms— Image Forensics, Forgery Localization, Decision Fusion, Background Information

1. INTRODUCTION

Nowadays it is increasingly common to find on the web digital images standing as proof of news or events. Despite their immediacy, digital images are also very easy to manipulate; for this reason, assessing their trustworthiness is of paramount importance. Authenticity verification for web images poses some challenging requirements. Firstly, it is desirable to precisely localize possibly forged regions, instead of just classifying the whole image as authentic or tampered. Secondly, it is common to have no knowledge about what kind of processing an image may have undergone, thus calling for the use of several complementary analysis tools. Finally, the huge amount of data on the web rules out supervised approaches, where a human patiently guides the image analysis.

While a good number of image authenticity verification tools have been developed in the last years, only a few of them

are able to localize manipulated regions in an unsupervised (i.e., fully automatic) way. Moreover, each tool usually works in very specific settings, thus requiring the analyst to combine the answer of multiple tools; doing this combination manually is time consuming and requires a strong expertise.

In this paper we propose a novel framework allowing to automatically merge the forgery localization maps produced by an arbitrary group of complementary tools. The proposed method, based on Dempster-Shafer Theory of Evidence (DST) [1], allows not only to fuse information coming from different unsupervised forensic tools, but also to exploit several kinds of background information to increase the reliability of the results. More precisely, our approach is able to exploit: i) *tool-based* information, since the fusion algorithm knows the reliability of each tool under different working conditions and exploits information about local and global properties of the analyzed content to better interpret the output of tools; ii) *trace-based* information, meaning that the fusion algorithm exploits knowledge of the compatibility relationships between traces and manages the case where two incompatible traces are simultaneously detected; iii) *semantic-based* information, which means exploiting the content of the analyzed image to improve the forgery localization map.

2. PREVIOUS WORKS ON UNSUPERVISED FORGERY LOCALIZATION

A first class of unsupervised forgery localization algorithms looks for the presence of tampered objects by decomposing the image under analysis into subparts, either using segmentation [2] or through a block-wise approach [3, 4]. However, since a sufficiently large portion of the image is usually needed for a reliable statistical analysis, only a coarse grained localization of tampering is possible with these methods.

Another class of algorithms allows automatic localization of the tampered regions with a fine-grained scale of $B \times B$ pixels (usually $B = 8$). The output of these methods is a likelihood map indicating for each pixel/block its probability

This work was partially supported by the European Union Seventh Framework Programme (FP7/2007-2013) under the MAVEN project, Grant Agreement 606058.

of being tampered. To the best of our knowledge, only few algorithms exploiting the presence of double JPEG compression [5–7] or the artifacts due to CFA interpolation [8] belong to this category. These approaches are strongly dependent on local and global properties of the image (content, dimension, compression etc) and often obtain noisy output maps.

An important limit of previous approaches is that they are based on the observation of a single forensic trace, whereas in practical scenarios, as those occurring on the web, the analysis of different footprints is needed. As to traces detected on the whole image, a number of techniques have been proposed to fuse the information at the feature level, i.e., by devising a complex classifier that accounts for multiple footprints [9,10]. Other approaches work at the score level, where the scalar output of the tools is considered during fusion [11, 12]. The overall performance of the above methods can be further improved by taking into account background information during fusion [13]. As to forgery localization, simple pixel-level fusion of different forensic tool outputs has been investigated in [14, 15]: however, these works do not consider tools reliability and compatibility, and they are based on very simple rules for fusion (sum, product, logical disjunction/conjunction).

3. ELEMENTS OF DEMPSTER-SHAFER THEORY

Dempster-Shafer Theory [1] is a mathematical theory providing a way to model uncertainty and to combine information coming from multiple sources. Let us denote with $\Theta = \{\theta_1, \theta_2, \dots, \theta_n\}$ the exhaustive set of mutually exclusive possible conclusions to be drawn. The frame of discernment of Θ is its power set 2^Θ , that is the set of all possible subsets of Θ . A Basic Belief Assignment (BBA) over Θ is a function $m^\Theta : 2^\Theta \rightarrow [0, 1]$ assigning a *mass* to elements of the frame of discernment associated to Θ , defined as:

$$m^\Theta(\emptyset) = 0; \quad \sum_{A \subseteq \Theta} m^\Theta(A) = 1 \quad (1)$$

where the summation is taken over all possible subset A of Θ . Intuitively, the mass assigned to a set is the amount of certainty supporting exactly that set, and not any of its subsets; for example it may be that $m^\Theta(\{\theta_1 \cup \theta_2\}) < m^\Theta(\{\theta_1\})$. The function accumulating the certainty about a set and all its subsets is called *belief* function:

Definition Given a BBA m^Θ over Θ , the Belief function $Bel : 2^\Theta \rightarrow [0, 1]$ is defined as follows:

$$Bel^\Theta(A) = \sum_{B \subseteq A} m^\Theta(B). \quad (2)$$

$Bel^\Theta(A)$ summarizes all our reasons to believe in A based on the available knowledge. Going back to the previous example, we surely have: $Bel^\Theta(\{\theta_1 \cup \theta_2\}) \geq Bel^\Theta(\{\theta_1\})$. The reader can find more details and properties in [1].

DST is widely known as a tool for combining the evidence coming from multiple independent sources of information. Indeed, given two BBAs m_1^Θ and m_2^Θ , we can obtain a fused BBA by using Dempster’s Combination Rule:

Definition Let Bel_1 and Bel_2 be belief functions over the same frame Θ with BBAs m_1 and m_2 . For all non-empty $X \subseteq \Theta$ the function m_{12} defined as:

$$m_{12}(X) = \frac{1}{1 - K} \cdot \sum_{\substack{A, B \subseteq \Theta: \\ A \cap B = X}} m_1(A)m_2(B) \quad (3)$$

where $K = \sum_{A, B: A \cap B = \emptyset} m_1(A)m_2(B)$, $K < 1$, is a BBA function defined over Θ and is called the *orthogonal sum* of Bel_1 and Bel_2 , denoted by $Bel_1 \oplus Bel_2$.

4. DST-BASED MULTI-CLUE ANALYSIS FOR FORGERY LOCALIZATION

The framework we propose aims at exploiting the output of an arbitrary set of unsupervised tamper localization algorithms and several kinds of background information so as to produce a single comprehensive and more reliable map.

Our system is reminiscent of the data fusion scheme described in [11]. In this scheme, the user manually selects a sufficiently large region and runs a set of tools assigning to the region a scalar value measuring the presence of a certain forensic trace in it. Then, the goal is to merge these outputs, by also taking into account some local properties of the region that may influence the reliability of the forensic tools. The way this is performed is briefly sketched below:

1. output from each tool is converted to a BBA about presence/absence of a trace in the selected region;
2. BBAs obtained from different tools are combined using Dempster’s rule (3), after applying belief extension for combining the information about different traces;
3. compatibility relationships between traces (modeled as a BBA) are introduced using Dempster’s rule;
4. final decision: the total belief that the region has been forged is computed based on the merged information.

The most intuitive approach to extend the above analysis to forgery localization would be to simply apply the whole procedure separately to each single element of the map (also called “analysis block”, from now on). However, this choice is potentially misleading because of the nature of forgery localization tools. Indeed, as stated in Section 1, the accuracy of forgery localization tools is strongly affected by the local properties of the image: for example, very smooth or saturated regions are critical for many tools (see, for example, [6,8]), so that values assumed by the map in those regions are less reliable. As a consequence, attention must be paid in properly interpreting the output of the tool locally. To this




Fig. 1. Block scheme of the proposed framework for forgery localization, where two tools A and B searching for a forensic trace α are considered. For the sake of clarity, global variables for Tool A are omitted in the drawing.

aim, for a forensic trace α , we define the set $\Theta_\alpha = \{t\alpha, n\alpha\}$, where $t\alpha$ is the proposition “trace α is present in the analysis block” and $n\alpha$ is the proposition “trace α is not present in the analysis block”. We model this *local* information provided by the tool τ with the following BBA over the frame Θ_α :

$$m_\tau^{\Theta_\alpha}(X) = \begin{cases} L_\tau(i) & \text{for } X = \{t\alpha\} \\ N_\tau(i) & \text{for } X = \{n\alpha\} \\ D_\tau(i) & \text{for } X = \{t\alpha\} \cup \{n\alpha\} \end{cases} . \quad (4)$$

In the above equation $L_\tau(i)$, $N_\tau(i)$ and $D_\tau(i)$ are scalar values obtained by interpreting the output of the tool in the i -th analysis block. It is here that *tool-based* background information enters the picture: besides considering the value of the localization map in the block i , a set of local properties of the image is evaluated (e.g., mean value or variance of pixels in the analysis block i) and used to determine the mentioned values for equation (4). To perform this mapping from tool outputs and background information to BBAs, we rely on the method recently proposed in [13]: such method exploits a set of training images to learn how local properties affect the output of the tool. Thus, given image and forgery localization map, using this approach we obtain values for (4) for each block of pixels. This stage of the framework is represented in the left-most side of Fig. 1 (“BBA mapping” blocks).

4.1. Global variables

Independently from the analysis domain (e.g., pixel or DCT domain), unsupervised forgery localization tools typically assume that the signal under analysis is the mixture of two components: one component deriving from parts of the image that were manipulated, and one deriving from unaltered parts [5, 6, 16]. A statistical model is defined for each component, and the parameters of the models are estimated from available data. Finally, each (block of) pixels is assigned a probability of belonging to each model, thus producing a forgery localization map. However, when for some reason the two components are not correctly separated, the produced localization map is practically useless, although it assigns a sensible value to each region. A simple way to understand whether the tool successfully separates the two components is to analyze the produced localization map as a whole: when the components are not separated, the whole map takes values in a narrow range, meaning that all pixels belong to the

same component, while the opposite happens when two components are separated.

The above discussion suggests that we cannot simply interpret elements of the map as “stand alone small blocks”, but we should also model the global information obtained from the map as a whole. We then introduce for each forensic trace α as reference, we define the frame $\Gamma_\alpha = \{T\alpha, N\alpha\}$ where $T\alpha$ is the proposition “the two components related to α were separated” while $N\alpha$ has the opposite meaning. After running a localization tool searching for α , a BBA over Γ_α must be defined. We are not forced to give a binary interpretation: indeed the border between the two cases is not always sharp. Hence, for a generic tool τ , we propose to model this information through the following BBA:

$$m_\tau^{\Gamma_\alpha}(X) = \begin{cases} (1 - W_\tau)G_\tau & \text{for } X = \{T\alpha\} \\ (1 - W_\tau)(1 - G_\tau) & \text{for } X = \{N\alpha\} \\ W_\tau & \text{for } X = \{T\alpha\} \cup \{N\alpha\} \end{cases} . \quad (5)$$

If the tool τ is based on model separation, then $G_\tau \in [0, 1]$ quantifies the confidence about the two components of the mixture being successfully separated, and $W_\tau = 0$. Instead, if τ is not based on model separation, we assign all the mass to the doubt by setting $W_\tau = 1$, thus yielding the neutral element for Dempster’s combination rule [1] and disabling the contribution of this BBA. This phase of the framework is drawn in the lower part of Fig. 1. Notice that, for the moment, the above BBA is not linked in any way to that in eq. (4) (they are also defined on different frames, Γ_α and Θ_α respectively), that is we are not still logically linking local and global information about the presence of the trace.

4.2. Inclusion of trace-based background information

Decision fusion is particularly interesting when the merged tools search for different traces, since, by knowing the theoretical properties of each forensic trace, in many cases the analyst can explicitly tell whether a combination of traces is plausible or not: this is what we call *trace-based* background information. As shown in [11], DST allows to write rather easily such information in terms of BBAs, and to combine it with the information provided by tools.

Also in this case, as we turn to forgery localization some noticeable differences appear. In the framework proposed

Θ_α	Γ_α	Θ_β	Γ_β	Plausible	Interpr.
t_α	T_α	n_β	T_β	Y	Tamp.
n_α	T_α	n_β	T_β	Y	Auth.
t_α	T_α	t_β	T_β	N	-

Table 1. Example of traces relationships.

in [11] each forensic trace is modelled with one variable, so that only relationships between different traces are to be considered. Here, instead, each trace is better represented with two variables (one referring to the local presence of the trace and one to the suitability of the global model). Hence, we also have a relationship between these two variables establishing the link between local and global information about the trace, and allowing to change the interpretation of the local output of the tool based on the global information. It is worth noting that the global information about the presence of one trace can also affect the interpretation of different forensic traces. Therefore, we write together these compatibility relationships through a table listing on rows the combinations of variables: each row is then labelled by the analyst as either plausible or not plausible. For plausible rows, the analyst also specifies the interpretation associated to that row in terms of authenticity of the block. Of course, this has to be done only once for a set of forensic traces. An example for two traces α and β is given in Table 1: the first row states that, for any image block where the global models of both trace α and β were successfully separated, it is plausible to find only the trace α and not the other; moreover, the interpretation associated to this combination is “the block is tampered”. The second row of the table tells that local absence of both traces is plausible and is to be interpreted as the block being authentic. The last row, instead, states that the two traces cannot be present simultaneously in the same block. The table is truncated for the sake of brevity; the complete version has 16 rows, even though it makes sense to write explicitly only plausible combinations.

Compatibility tables can be easily written in terms of a BBA as follows: for a given set \mathcal{T} of considered traces, let us define as $\Psi = \prod_{j \in \mathcal{T}} \Theta_j \times \Gamma_j$ the common frame of discernment, where \prod and \times denote the Cartesian product. Let us also denote by $\Psi_{\text{PL}} \subseteq \Psi$ the union set of all combinations that are considered plausible. Then, the following BBA declares that implausible combinations have to be considered as conflicting information (in Fig. 1 it corresponds to the block outputting m_{rel}^Ψ):

$$m_{\text{rel}}^\Psi(X) = \begin{cases} 1 & \text{for } X \in \Psi_{\text{PL}} \\ 0 & \text{for } X \notin \Psi_{\text{PL}} \end{cases} \quad (6)$$

4.3. Obtaining the fused localization map

By applying Dempster’s combination rule to the BBA resulting from traces relationship and those available from single tools, we obtain a single BBA summarizing the available information. Then, it makes sense to compute the belief of the

set composed by all plausible combinations whose interpretation is “tampered”, using equation (2). Notice that this formula is computed only once for a given set of forensic traces, then it can be stored and evaluated when needed in $O(N)$ time. By evaluating the formula for each analysis block of an image, a map taking values in $[0,1]$ is produced, which tells the total belief for each block of being tampered.

As localization tools process each analysis block independently of the others [5,6,8], the resulting localization maps are typically affected by noise. In some cases, authors proposed to filter the map (as with median filtering [6]), but this solution could be not sufficient when several maps have to be fused. Moreover, the use of filtering based on fixed window rises the problem of how to set the window size. As solution, we propose to exploit what we call *semantic-based* background information, meaning that we let the content of the analyzed image to drive the map processing. Recently, [17] proposed to use guided filtering [18] to accomplish this task: guided filter computes the output by considering the content of the guidance image. In our case, the input is the localization map and the guidance image is the image under inspection; moreover, the guided filter transfers the structures of the guidance image to the filtered output in $O(N)$ time.

5. EXPERIMENTAL RESULTS

In this section we discuss the experiments that we carried out. The tools we employ are based on *aligned* double JPEG compression (AJPEG) footprints [6], *non-aligned* double JPEG (NAJPEG) footprints [16] and Color Filter Array (CFA) inconsistencies [8]. Let us summarize the underlying scenarios. In [6] the case where an original JPEG image, after some localized forgery, is saved again in JPEG format is considered: DCT coefficients of unmodified areas undergo a double *aligned* JPEG compression, thus exhibiting double quantization (DQ) artifacts, while DCT coefficients of forged areas do not. The tool in [16] assumes that a region from a JPEG image is pasted onto a host image that does not exhibit the same JPEG compression statistics, and that the resulting image is JPEG re-compressed: here, the forged region likely exhibits *not aligned* double compression artifacts, whereas the original region does not. In the scenario used in [8], a local forgery destroys the correlation introduced by in-camera *demosaicing*: the forged region does not show CFA artifacts, whereas the remaining part of the image does.

In order to define the mapping from the localization maps to BBAs (Eq. (4)), we adopt the method proposed in [13], choosing the following set of properties to locally characterize the reliability of each tool τ : the value of the last compression factor (if any) q_2 , the mean value μ and the std. deviation σ of pixels intensity and the first compression factor q_1 , if present. Note that q_1 is not directly observable, but it is estimated by AJPEG and NAJPEG tools, and it is employed only for CFA traces of CFA artifacts could be removed by strong

past compression [8]. The generic analysis block is thus described by the vector $v = (o_\tau, q_2, \mu, \sigma, q_1)$, where o_τ denotes the value of the block in the map produced by tool τ (in our case, $\tau \in \mathcal{T} = \{\text{AJPEG}; \text{NAJPEG}; \text{CFA}\}$). By applying the approach proposed in [13], each vector is associated to scalar values L_τ , N_τ and D_τ (see Eq. 4); as to the parameters required in [13], we used $\alpha = 0.85$ and $k = 12$ for all tools, whereas $\gamma = 0.5$ for CFA tool, $\gamma = 512$ for AJPEG tool and $\gamma = 2048$ for NAJPEG tool. Values were obtained through 5-fold cross validation and grid search.

Finally, we define an empirical method to assign values to global variables, telling to what extent the tool successfully separated the two components for its own trace. Since the considered tools are based on model separation, according to eq. (5) we set $W_\tau = 0 \forall \tau \in \mathcal{T}$, and we define a linear piecewise function: $G_\tau(\rho) = \rho/a$ for $\rho < a$, and $G_\tau(\rho) = 1$ for $\rho \geq a$, where ρ is the percentage of blocks belonging to the less populated model, as explained in Sec. 4.1. The value of a , representing the minimum percentage of blocks allowing a model to be detected, was set from experimental evidence to $a = 1/8$. So, two components are separated if at least 1/8 of the blocks shows the footprints searched for.

5.1. Results

Here we show the improvements in localizing forgeries in an unsupervised scenario. To quantify it, we generate three different sets of images to train and test the proposed framework. Firstly, we define a *training* set to train the BBA mapping module, incorporating *tool-based* background information. The second step is to design a *testing* dataset to compare the performance of each tool employed individually with respect to those of the framework. It is worth noting that we assume a *blind* case, i.e. each tool is applied without any a priori information about the type of tampering. Finally, we built a dataset of realistic spliced images to show the capabilities of localizing a forged region, listed below.

Training: Starting from 100 1024×1024 uncompressed TIFF images, three different tampering are applied separately, in such a way that the traces detected by each algorithm have been inserted (or deleted) from the left half of each image. For the AJPEG and NAJPEG traces, the quality factors of the first and second compression are in $\{50, 60, 70, 80, 90, 100\}$, whereas for the CFA footprint, the quality factors employed are in $\{50, 60, 70, 80, 90, 100, \text{Inf}\}$, where Inf represents the case of TIFF uncompressed images. By combining all possible compression factors, we obtain a set composed by 3600 images for AJPEG, 3600 for NAJPEG and 700 for CFA case.

Testing: Starting from 50 uncompressed TIFF images, different from the training set, we apply the same tampering as before to the central 512×512 block. For AJPEG and NAJPEG traces, the quality factors of the first compression are in $\{60, 70\}$, whereas the quality factors of the second are in $\{80, 90\}$. For the CFA based tampering, a median filtering is

	Unweighed	Weighed
AJPEG	0.854	-
NAJPEG	0.607	-
CFA	0.588	-
Sum	0.681	0.627
Product	0.686	0.556
DST	0.692	0.895

Table 2. AUC values of different localization methods based on single tool (AJPEG, NAJPEG, CFA), or fused (Sum, Product and DST) by using or not global variable based weighting.

applied to remove traces of CFA artifacts. Overall, 750 test images have been created: 200 with AJPEG tampering, 200 with a NAJPEG tampering, 150 with CFA tampering and 200 containing AJPEG and NAJPEG traces at the same time.

Realistic: 19 realistic forgeries have been created through a *cut and past* strategy, by inserting a content (i.e. an object) from an image onto another one, and keeping track of the forgery position. The set is composed of 4 TIFF images, whereby an object (without CFA artifacts) is pasted onto another (with CFA artifacts), 6 images with AJPEG footprints, 5 images with NAJPEG footprints and 4 images whereby objects with NAJPEG traces have been inserted in images with AJPEG traces. All forgeries were made in such a way that each footprint is easily detected, since the aim of this dataset is to evaluate the capability of localizing a realistic forgery.

To prove the validity of the framework, we use the *true positive rate* (R_{TP}), measuring the fraction of tampered blocks correctly detected as forgery, and the *false positive rate* (R_{FP}), measuring the fraction of unchanged blocks erroneously detected as forgery. The overall performance of the compared methods are evaluated by plotting its *receiver operating characteristic* (ROC) curve. The *area under the curve* (AUC) is finally employed to summarize the discrimination capability of detectors. The first test is carried out on the *testing* dataset, aiming to compare our framework to each tool applied independently and in a blind way, and the methods proposed in [14, 15]. The performance, evaluated in terms of AUC, show that the DST-based framework outperforms the single detectors, as shown in Table 2. It is worth noting that no post-filtering has been applied. As we can see, the proposed framework has the best capability of localizing forgeries, and the introduction of global variables dramatically impacts the performance, since their use provides further information about the reliability of the value given by a tool.

Finally, we present the localization capability of the framework when applied to the *realistic* dataset. In Fig. 2, we show the performance of the method without post-filtering and with guided filtering. Moreover, a comparison with each tool performance is proposed. As expected, the refinement by using guided filtering increases the accuracy in localizing realistic forgeries. Even in this case, the DST-based framework has better capabilities with respect to each single tool.




Fig. 2. Localization capability without post-filtering (dotted curve), with guided filtering (solid) and of each single tool AJPEG (circled), NAJPEG (crossed) and CFA (dashed).

6. CONCLUSIONS AND FUTURE WORK

In this paper a framework for unsupervised multi-clue forgery localization has been proposed, which merges information provided by a set of forensic tools with background information freely available to the analyst. Such a framework exploits the peculiar properties of those localization tools that are based on mixture models, by introducing global variables that are taken into account by the system. Although the way we assigned values to such variables is still rather empirical, their impact on the overall performance is dramatic. The formalization of global variable assignments and the extension to the case of copy-move detectors, that cannot distinguish between original and pasted regions, is left for future work.

7. REFERENCES

- [1] G. Shafer, *A Mathematical Theory of Evidence*, Princeton University Press, 1976.
- [2] M. Barni, A. Costanzo, and L. Sabatini, "Identification of cut & paste tampering by means of double-JPEG detection and image segmentation," in *IEEE ISCS*, 2010, pp. 1687–1690.
- [3] M. Chen, J. Fridrich, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE TIFS*, vol. 3, no. 1, pp. 74–90, 2008.
- [4] A. Swaminathan, M. Wu, and K. J. R. Liu, "Digital image forensics via intrinsic fingerprints," *IEEE TIFS*, vol. 3, no. 1, pp. 101–117, 2008.
- [5] Z. C. Lin, J. F. He, X. Tang, and C. K. Tang, "Fast, automatic and fine-grained tampered JPEG image detection via DCT coefficient analysis," *Pattern Recognition*, vol. 42, no. 11, pp. 2492–2501, 2009.
- [6] T. Bianchi, A. De Rosa, and A. Piva, "Improved DCT coefficient analysis for forgery localization in JPEG images," in *IEEE ICASSP*, 2011, pp. 2444–2447.
- [7] T. Bianchi and A. Piva, "Image forgery localization via block-grained analysis of jpeg artifacts," *IEEE TIFS*, vol. 7, no. 3, pp. 1003–1017, 2012.
- [8] P. Ferrara, T. Bianchi, A. De Rosa, and A. Piva, "Image forgery localization via fine-grained analysis of cfa artifacts," *IEEE TIFS*, vol. 7, no. 5, pp. 1566–1577, 2012.
- [9] Y.-F. Hsu and S.-F. Chang, "Statistical fusion of multiple cues for image tampering detection," in *Conf. on Signals, Systems and Computers*, 2008, pp. 1386–1390.
- [10] P. Zhang and X. Kong, "Detecting image tampering using feature fusion," in *ARES*, 2009, pp. 335–340.
- [11] M. Fontani, T. Bianchi, A. De Rosa, A. Piva, and M. Barni, "A framework for decision fusion in image forensics based on Dempster-Shafer Theory of Evidence," *IEEE TIFS*, vol. 8, no. 4, pp. 593–607, 2013.
- [12] M. Barni and A. Costanzo, "A fuzzy approach to deal with uncertainty in image forensics," *Signal Processing: Image Communication*, vol. 27, no. 9, pp. 998–1010, 2012.
- [13] M. Fontani, E. Argones-Rua, C. Troncoso, and M. Barni, "The watchful forensic analyst: Multi-clue information fusion with background knowledge," in *IEEE WIFS*, 2013, pp. 120–125.
- [14] D. Cozzolino, F. Gargiulo, C. Sansone, and L. Verdoliva, "Multiple classifier systems for image forgery detection," in *Image Analysis and Processing*, vol. 8157, pp. 259–268. Springer, 2013.
- [15] L. Gaborini, P. Bestagini, S. Milani, M. Tagliasacchi, and S. Tubaro, "Multi-clue image tampering localization," in *IEEE WIFS*, 2014, pp. 1–6.
- [16] Tiziano Bianchi and Alessandro Piva, "Detection of non-aligned double JPEG compression with estimation of primary compression parameters," in *IEEE ICIP*, 2011, pp. 1929–1932.
- [17] G. Chierchia, D. Cozzolino, G. Poggi, C. Sansone, and L. Verdoliva, "Guided filtering for PRNU-based localization of small-size image forgeries," in *Proc. of IEEE ICASSP*, 2014.
- [18] K. He, J. Sun, and X. Tang, "Guided image filtering," in *Computer Vision–ECCV*, pp. 1–14. Springer, 2010.