

Detecting morphological filtering of binary images

Francesco G.B. De Natale, Giulia Boato

Abstract—Morphological operators are widely used in binary image processing for several purposes, such as removing noise, detecting contours or particular structures, regularizing shapes. In particular, morphological filters are largely adopted in scanned documents to correct the artifacts caused by acquisition and binarization, as well as other processing. In this paper we propose a novel approach for forensics detection of morphological filtering on binary images. The proposed technique exploits some mathematical properties of the two basic morphologic operators, erosion and dilation, to define an algorithm able not only to detect the application of the filter, but also to estimate the shape of the relevant structuring element. Experimental tests demonstrate that the technique is effective and robust to the most common operations performed on binary image documents.

Index Terms—Image Forensics, Mathematical morphology, Filtering detection.

I. INTRODUCTION

The detection of forgeries in digital documents is becoming a key issue in image forensics. The possibility of easily manipulating and reproducing paper documents through scanning, editing and printing, facilitates a number of frauds and can also be seen as an accessory activity in various criminal actions. Most typical abuses are the generation of false documents and the tampering of existing ones, with the purpose of giving official form to deceptive information, stealing signature and identity, reproducing copyrighted material, altering sensitive written information such as correspondence, certificates, etc.

In this context, it is essential to provide forensics examiners with instruments that can unveil the most frequent manipulations of digital documents. Such manipulations include cut&paste and other composition techniques, textual modifications, printing/acquisition sequences, digital filtering. In particular, filters are often used as a post-processing operation to remove artifacts and to conceal the traces associated to previous operations and are therefore considered a hint of possible malicious modifications.

Most part of the binary image forensics technologies so far proposed are concerned with the identification of scanning and printing devices used to acquire, reproduce and possibly counterfeit documents. In [1] the authors provide a thorough description of the problems related to printed document trustworthiness, and review the most significant state of the art methodologies for scanner and printer forensics, mainly focusing on the generation of fallacious documents and the tampering of existing ones. Additional information in this area can be found in [2][3], and references thereby. In particular, in [4] a detailed survey throughout photo forensics is provided

with particular attention to all different types of analyses relying on physics, geometry, optic, file or pixel details.

As far as filtering is concerned, binary images allow a limited range of processing operations, due to the pixel encoding limitations. Typical operators used in binary image filtering have non-linear nature and include rank-order, majority, and morphological filters. The latter, in particular, have gained great popularity thanks to the ease of implementation, fast processing, and adaptability to different purposes. A lot of work has been spent in recent years on the blind detection of non-linear filtering of graylevel images [5][6][7], and in particular the median filter [8][9][10][11][12][13][14]. Part of these methods may be adapted to work on binary images upon suitable modification, although they are not specifically targeted to the characteristics of such images.

On the contrary, despite the widespread use of morphological filtering in binary image processing, little or no effort has been spent as of today on the development of effective forensics detection methodologies able to reveal their application. In the last years, a few researchers made use of morphology as a tool to implement effective forensic detectors of image manipulations. In [15] for instance, mathematical morphology was adopted as a component of a splicing detector, while in [16] morphological operators are used to reveal the size of the blocks used in image compression. Nevertheless, to the best of our knowledge, no attempt has been made to produce forensics approaches to blindly detect the application of any kind of morphological filter to a digital image.

In this paper we address for the first time the above problem, and in particular we propose an algorithm able to automatically detect the traces of the application of a morphological erosion or dilation to a binary image. To this purpose, we review some basic principles of mathematical morphology, and we derive some related properties that lead to the theoretical formulation of a detector. Then, we define a forensics algorithm that, scanning a minimal number of configurations, allows detecting not only the application and type of the filter, but also the shape of the structuring element used. In the same line of the work in [17] on median filtering detection, the proposed detector does not make use of statistical approaches or machine learning technologies to take a decision, but utilizes some properties of the filtered signal that allow distinguishing it from an original one in a deterministic way.

The rest of the paper is organized as follows: in Section II the proposed method is described in detail, starting from the mathematical formulation to derive the algorithmic structure of the detector; in Section III we introduce the dataset used for the testing and the experimental setup, and we present and discuss the results achieved. Finally, in Section IV we draw the conclusions of our work and we highlight some possible future directions of our research.

II. PROPOSED APPROACH

In this section, we first introduce the theoretical foundations of the problem, and then derive from it the proposed detection algorithm.

A. Review of mathematical morphology theory, and foundations of the proposed detector

Mathematical morphology defines a class of nonlinear filters first introduced in the early 60s by G. Matheron and J. Serra in the framework of their mineralogic studies [18] (see also [19] for a comprehensive reading). All the morphological filters originate from different combinations of the two fundamental operators, *erosion* and *dilation*, and a kernel called *structuring element*, defined as a binary mask with a given shape and a reference point. The shape of the structuring element determines the effect of the filter on the image. While the original theory was thought for binary images, mathematical morphology was later generalized to the case of grayscale pictures [20].

Focusing on the binary case, given an image A (the *object*) and a structuring element B , the two basic operators, erosion and dilation, respectively, are defined as follows:

$$A \ominus B = \{x \text{ s.t. } B_x \subseteq A\} \quad (1)$$

$$A \oplus B = \{x \text{ s.t. } B_x^S \cap A \neq \emptyset\} \quad (2)$$

where B_x represents the structuring element B with the reference point positioned in x , while B_x^S represents the reflective rotation of B positioned in x . Figure 1 shows an example of erosion and dilation of a character with the depicted structuring element.

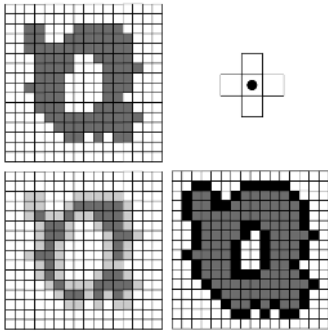


Fig. 1: Example of erosion (bottom left) and dilation (bottom right) of a character (top left) with a 'cross' structuring element (top right). Light gray pixels represent the points canceled by the erosion, while black pixels represent the points added by the dilation.

We can then define a number of composite morphological filters as a combination of the above. The most common composite operators are the *open* and *close* operators, defined as the sequence erosion-dilation and dilation-erosion, respectively:

$$A \circ B = A \ominus B \oplus B \quad (3)$$

$$A \bullet B = A \oplus B \ominus B. \quad (4)$$

Other classical filters include *hit-or-miss* and *skeletonization* operators, largely used in shape analysis to locate corners or isolated points, and to represent shapes.

Morphological operators enjoy a number of well-known properties, including the following, which will be used later in this work:

- (i) translation invariance: the translation of the filtered image depends on the position of the reference point only
- (ii) commutativity of the dilation, $A \oplus B = B \oplus A$
- (iii) associativity: the cascade of multiple filters can be transformed into a single filter whose structuring element is the morphological combination of the original elements, according to the following equations:

$$A \ominus B_1 \ominus B_2 = A \ominus \{B_1 \oplus B_2\} \quad (5)$$

$$A \oplus B_1 \oplus B_2 = A \oplus \{B_1 \oplus B_2\} \quad (6)$$

- (iv) idempotence of open and close: the iteration of open and close operators with the same structuring element does not produce any further modification of the image:

$$A \circ B \circ B = A \circ B \quad (7)$$

$$A \bullet B \bullet B = A \bullet B. \quad (8)$$

Although (iv) is the most known property of open and close operators, in our work we will better use another property, which establishes a similar equality without requiring a complete sequence of open or close operations, namely:

Theorem 1: Let $C = A \ominus B$, then $C \bullet B = C$. Equivalently, let $C = A \oplus B$, then $C \circ B = C$.

For a demonstration of Theorem 1, please look at [21], p. 16. The statement of Theorem 1 can be rewritten as follows:

$$A \ominus B = A \ominus B \oplus B \ominus B \quad (9)$$

and equivalently

$$A \oplus B = A \oplus B \ominus B \oplus B. \quad (10)$$

Now we can prove the following:

Theorem 2: Let $C = A \ominus B$, then $\forall D$ such that $\exists E \mid D \oplus E = B$ we have that $C \bullet D = C$. Mutatis mutandis, given $C = A \oplus B$, then $C \circ D = C$.

Proof: $C = A \ominus B = A \ominus D \oplus E = A \ominus E \oplus D = A \ominus E \oplus D$ where the third equality holds for property (ii), and the fourth for property (iii). Now, given $F = A \ominus E$, for Theorem 1 we have: $C = F \oplus D = \{F \oplus D\} \bullet D = F \oplus D \oplus D \ominus D = A \ominus E \oplus D \oplus D \ominus D = C \oplus D \ominus D = C \bullet D$. \square

Theorem 2 states that the equality established by Theorem 1 holds also if the open or close operations are performed with any structuring element D that can be dilated to obtain B .

In the following subsection we will exploit the results of the two above theorems to derive a forensic tool able to detect the application of morphological filters to binary images. In particular, we will see that Theorem 1 provides a practical method to detect the last erosion/dilation operation performed on an image, while Theorem 2 allows significantly reducing the computational complexity of the detector, limiting the number of structuring elements to be considered in the detection process.

B. Detection algorithm

As stated in the previous section, basically all mathematical morphology operators can be defined in terms of combinations of erosion and dilation filters along with set operators such as intersection and union. For this reason, in this work we focus on these two fundamental operators, aiming at detecting their application to binary images.

Figure 2 shows a block diagram of the proposed detector. The idea is the following: given an unknown binary image I , we want to determine whether it has been filtered with a morphological erosion, a morphological dilation, or none. Let suppose that I is the result of an erosion operator with a structuring element B , then we will have $I = I' \ominus B$, where I' is in general not available. Given Theorem 1, and assuming that we know B , we can state that:

$$I \bullet B = I \oplus B \ominus B = I' \ominus B \oplus B \ominus B = I' \ominus B = I. \quad (11)$$

This provides a practical way to check if I is the result of an erosion with B : it will be sufficient to apply to I a sequence of dilation and erosion with B and check if the result is pixelwise equal to I . Reverting erosion and dilation operators in Eq. (11) provides the relevant dilation detector. Clearly, if the erosion detector is applied to a dilated or to a non-filtered image, the result will in general differ from I , and the same happens if the detection is done with a structuring element B' different from B .




Fig. 2: Block diagram of the detector.

It is to be observed that the assumption of knowing B is obviously unrealistic in general. This implies that the detector will have to scan a set of masks in order to find a possible match. In this work we are mainly interested in the use of morphology as a regularization tool, such as removing noise and artifacts. This is clearly the most interesting case from the point of view of image forensics, as those filters do not change

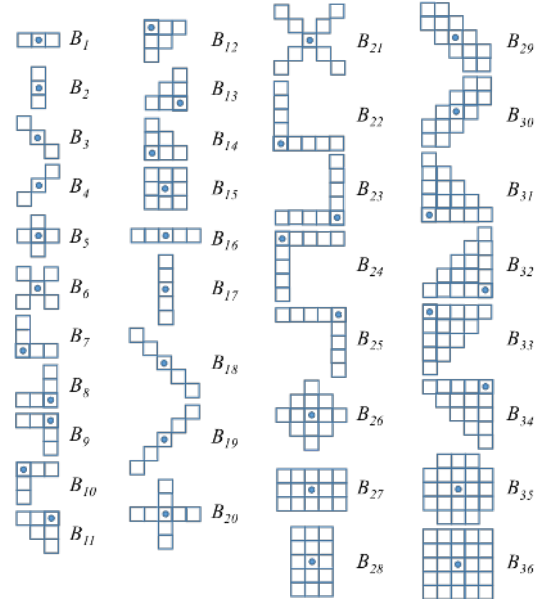


Fig. 3: Set Ω of structuring elements used for simulations.

the nature of the document. In fact, the use of morphology as a detector (e.g., corners, skeleton, shape analysis) typically produces an image completely different from the original. Accordingly, we considered a set of masks $\Omega = \{B_i\}$ for $i = 1, \dots, N$ including all the filters that are commonly used in the literature for shape regularization, and are characterized by high symmetry, also with directional properties. Figure 3 reports the set Ω used in this paper, which includes $N = 36$ structuring elements, with a maximum size of 5×5 pixels. In each mask the reference point has been placed in the symmetry center and indicated by a dot: possible shifts of the reference point will simply produce a translation of the output image, which is irrelevant for our analysis thanks to property (i).

Since binary erosion and dilation are computationally effective operations, it could be acceptable to use a brute-force approach, applying the detector to the whole set of masks. Nevertheless, Theorem 2 allows a significant optimization. In fact, we can easily reverse the implications of Theorem 2 to derive the following statement: if the test with a structuring element B_i fails, every other structuring element B_j that can be obtained by dilating B_i in any possible way will fail as well. Thus, for each element B_i we can define a set $\hat{\Omega}_i$ made of all the masks $B_j \in \Omega$ that can be obtained by dilating B_i . If a no-match is found for B_i , all the elements in $\hat{\Omega}_i$ (also called *implications* in the following) can be excluded from further analysis thanks to Theorem 2. The implications of each element B_i can be calculated a-priori, thus reducing the solution space and the number of tests to be performed. Table I reports the list of implications for the elements B_i depicted in Figure 3. For the sake of clarity, the dilation operations that allow excluding each specific $B_j \in \hat{\Omega}_i$ are also specified in parenthesis. It is to be pointed out that the elements in $\hat{\Omega}_i$ could also be generated by dilating B_i with elements not included in Ω . This is the case, e.g., of elements B'_1, B'_2 appearing in the dilation of masks B_3 and B_4 , respectively.

B_i	$\hat{\Omega}_i$	
B_1	$B_{15} (B_1 \oplus B_2)$	$B_{16} (B_1 \oplus B_1)$
	$B_{27} (B_1 \oplus B_{15})$	$B_{28} (B_1 \oplus B_{17})$
	$B_{35} (B_1 \oplus B_2 \oplus B_5)$	$B_{36} (B_1 \oplus B_{28})$
B_2	$B_{15} (B_2 \oplus B_1)$	$B_{17} (B_2 \oplus B_2)$
	$B_{27} (B_2 \oplus B_{16})$	$B_{28} (B_2 \oplus B_{15})$
	$B_{35} (B_2 \oplus B_1 \oplus B_5)$	$B_{36} (B_2 \oplus B_{27})$
B_3	$B_{18} (B_3 \oplus B_3)$	$B_{29} (B_3 \oplus B'_1)$
B_4	$B_{19} (B_4 \oplus B_4)$	$B_{30} (B_4 \oplus B'_2)$
B_5	$B_{26} (B_5 \oplus B_5)$	$B_{35} (B_5 \oplus B_{15})$
B_6	$B_{36} (B_6 \oplus B_{15})$	
B_7	$B_{31} (B_7 \oplus B_{14})$	
B_8	$B_{32} (B_8 \oplus B_{13})$	
B_9	$B_{34} (B_9 \oplus B_{11})$	
B_{10}	$B_{33} (B_{10} \oplus B_{12})$	
B_{11}	$B_{34} (B_{11} \oplus B_9)$	
B_{12}	$B_{33} (B_{12} \oplus B_{10})$	
B_{13}	$B_{32} (B_{13} \oplus B_8)$	
B_{14}	$B_{31} (B_{14} \oplus B_7)$	
B_{15}	$B_{27} (B_{15} \oplus B_1)$	$B_{28} (B_{15} \oplus B_2)$
	$B_{35} (B_{15} \oplus B_5)$	$B_{36} (B_{15} \oplus B_{15})$
B_{16}	$B_{27} (B_{16} \oplus B_2)$	$B_{36} (B_{16} \oplus B_{17})$
	$B_{28} (B_{16} \oplus B_1)$	$B_{36} (B_{16} \oplus B_{16})$
B_{27}	$B_{36} (B_{27} \oplus B_2)$	
B_{28}	$B_{36} (B_{28} \oplus B_1)$	

TABLE I: Structuring elements exclusions implied by Theorem 2: for each B_i the corresponding set $\hat{\Omega}_i$ is reported, along with the generating morphological operations (in parenthesis).

Algorithm 1 Erosion detector

Input: Test image I , structuring elements $\Omega = \{B_i\}, i \in [1, N]$

Output: Positive or negative detection of erosion on I

Method:

- 1: $\hat{\Omega} = \emptyset, E = \emptyset$
 - 2: **for** $i = 1, \dots, N$ **do**
 - 3: **if** $B_i \in \Omega \setminus \hat{\Omega}$ **then**
 - 4: compute $I_C \bullet B_i$
 - 5: **if** $I_C = I$ **then** $E = E \cup B_i$ // match
 - 6: **else** $\hat{\Omega} = \hat{\Omega} \cup \hat{\Omega}_i$ // no match, exclude implications
 - 7: **end for**
 - 8: **if** $E == \emptyset$ **then return** negative
 - 9: **else return** positive to mask $B_{max}, max = \max i$ s.t. $B_i \in E$
-

Accordingly, we can arrange the sequence of checks starting from smaller (thus, faster) filters, and progressively excluding the larger elements that fulfill the above rules. For instance, if masks B_1, B_2, B_3, B_4 fail, it is possible to exclude from the check $\hat{\Omega}_1 \cup \hat{\Omega}_2 \cup \hat{\Omega}_3 \cup \hat{\Omega}_4 = \{B_{15}, B_{16}, B_{17}, B_{18}, B_{19}, B_{27}, B_{28}, B_{29}, B_{30}, B_{35}, B_{36}\}$. Similarly, if the first 10 small structuring elements fail, we can exclude 16 bigger masks from the test. This allows us to reduce significantly the computation.

On the contrary, if an element B_i succeeds, this cannot exclude that a larger element $B_j \in \hat{\Omega}_i$ could succeed as well. If this happens, the detector will output the largest structuring element in the positive subset. The procedure is summarized in Algorithm 1 for the erosion detector. The dilation detector works exactly the same way, by only substituting close with open in line 4 (equation (11)).

It is to be pointed out that the above property excludes the possibility for the detector to determine multiple sequential filtering, in the sense that the detector will be unable to

distinguish between the sequence itself and its concatenation into a single combined filter. For instance, if a dilation with mask B_{27} is detected, this may also result from a sequential dilation with masks B_2 - B_{16} (in either order), B_1 - B_{15} , B_1 - B_2 - B_1 , and so on, according to the implications of Table I.

III. TESTING AND EXPERIMENTAL RESULTS

In this Section we present the experimental validation of the proposed detector. First, we introduce the datasets used in the experiments, then we analyze the results in terms of accuracy and robustness, finally we evaluate the possible application of the tool to detect splicing operations.

A. Datasets and experimental setup

The proposed detector was tested on two public datasets with different characteristics, both associated to typical applications of binary documents.

The first dataset is related to the archiving of scanned documents, including text and graphics. The LRDE Document Binarization Dataset¹ is an open dataset of scanned newspaper pages, originally thought for validation of binarization algorithms. In our case, we simply binarized each image by converting it to 8 bit grayscale, applying a fixed threshold at 128, and re-encoding at 1 bpp. The resulting images have a fixed size of 2516×3272 pixels, corresponding to a resolution of 300 dpi. The number of images used in our experiments is 125. For more details about LRDE dataset, please look at [22]. In the following we will refer to this dataset as LRDE. Figure 4 shows an example of a typical LRDE page.

The second dataset is related to handwritten documents. To this purpose, we selected the IAM handwriting database [23]. We randomly chose 125 out of the 1539 scanned page documents, each one with a resolution of 2479×3542 pixels at 300 dpi, with a depth of 8 bits, grayscale. Also in this case each image was binarized with a fixed threshold of 128 and re-encoded at 1 bpp. In the following we will refer to this dataset as IAM. Figure 5 shows an example of a typical IAM page.

For both dataset, different settings of the threshold used for binarization have been experimentally proven to be influential in our analysis, as far as they provide a readable binary document. Accordingly, a fixed threshold has been adopted to ease the replicability of the experiments.

It is to be pointed out that, given the nature of the detector, the probability of false negatives in case of no further processing is zero by definition. In fact, if the image is filtered and no other operation has been performed on it, the theoretical formulation guarantees that the counter of the differences at the output returns zero. Vice versa, there is a non null probability of false positives, which are typically related to input images showing very limited or null content. As a matter of fact, given the extreme case of an input image containing no objects, the effect of the filter is null and therefore cannot be detected. A similar situation may rarely happen also when the

¹Copyright (c) 2012. EPITA Research and Development Laboratory (LRDE) with permission from Le Nouvel Observateur.

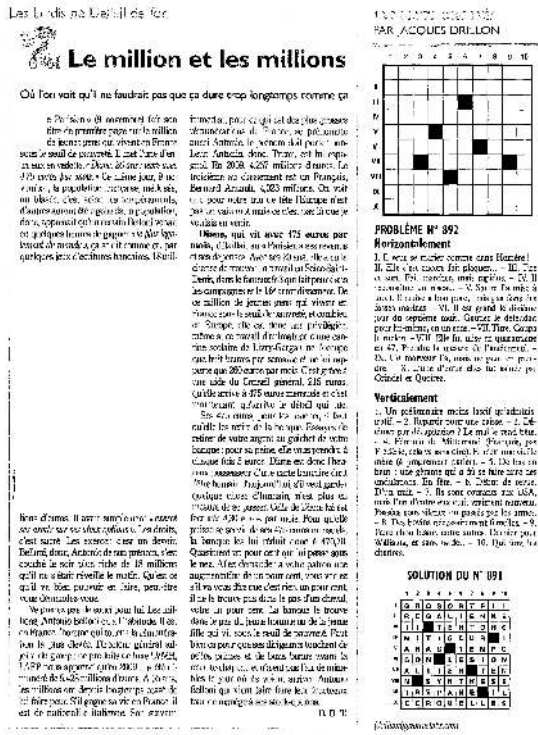


Fig. 4: A typical example of LRDE scanned image.

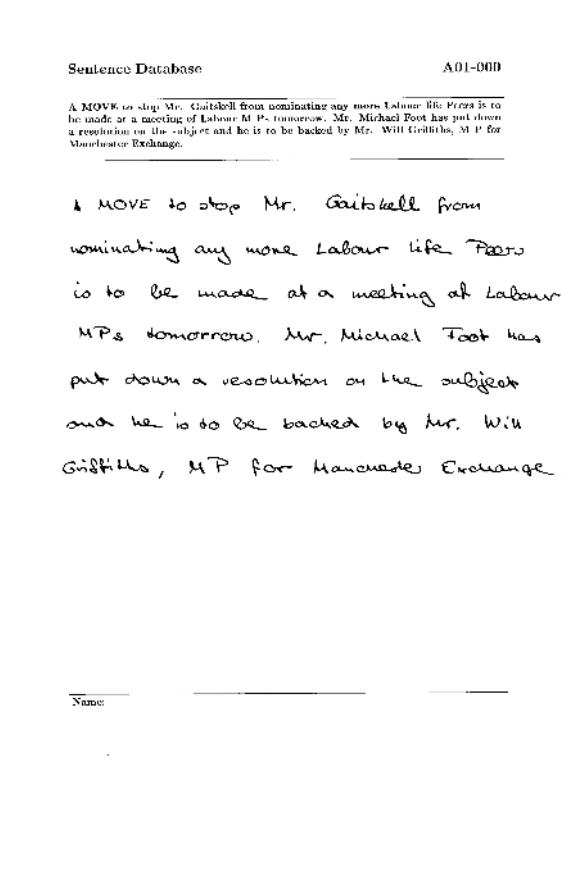


Fig. 5: A typical example of IAM scanned image.

shapes contained in the image are very regular, then being immune to open and close operators. Accordingly, in the first set of experiments we will present the accuracy of the detector in terms of false positives rate only, both for the full image and for smaller blocks. This last analysis will also enable revealing localized manipulation and will be used in the following splicing detection analysis.

A final consideration has to be made about comparisons. As already stated, to the best of our knowledge our method is the very first attempt to detect morphological image filtering. As a consequence, there are no state-of-art approaches to be considered for comparison purposes. Nevertheless, we will show that the proposed method provides nearly 100% accuracy in the absence of attacks, and is robust to large part of the common post-processing operations typically performed on binary images, making it a significant benchmark for the solution of this problem.

B. Experimental results and discussion

The first set of tests aims at assessing the accuracy of the detector on the two datasets in the absence of post-filtering processing. To this purpose, each image was filtered with erosion and dilation operators, using all the structuring elements shown in Figure 3, thus obtaining 9000 (125 × 36 × 2) filtered images for each dataset. Every filtered image, as well as the original ones, were then input to both erosion and dilation detectors. Each detector was run with all the probe masks, according to the algorithm described in Section II-B. As already explained, the detector returns the largest inclusive structuring element that outputs a zero value, or a non-detect condition if no one mask returns a null value. The experimental tests confirmed the theoretical assumption that false negatives are null, thus shifting the attention to false positives, i.e., non filtered images detected as filtered.

The first important result is that both detectors provided a 100% accuracy over both datasets, being able to correctly classify each full image as eroded, dilated or original (see Table II).

FILTER	DETECTOR	
	erosion	dilation
no filter	0%	0%
erosion	100%	0%
dilation	0%	100%

TABLE II: Detection accuracy on both datasets.

In order to assess the accuracy of the detector in localizing the filtered area also on small patches, we performed a further test where we applied the detector to image blocks of variable size. To this purpose, we randomly selected a large number of non-null blocks scanning a subset of images from the dataset, with a dimension of 64 × 64, 128 × 128, and 256 × 256, respectively (corresponding to a patch size from 0.2" to 0.85" at the given resolution). The cardinality of the resulting analysed blocks after filtering were 72373, 16269, and 4122 for erosion, and 181027, 50226, and 12792 for dilation, due

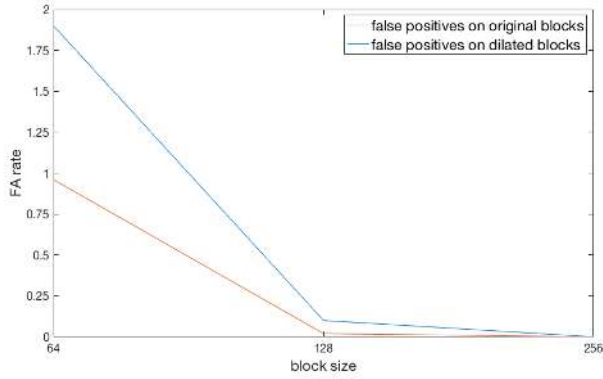


Fig. 6: Erosion detector false alarm (FA) rate for different block sizes.

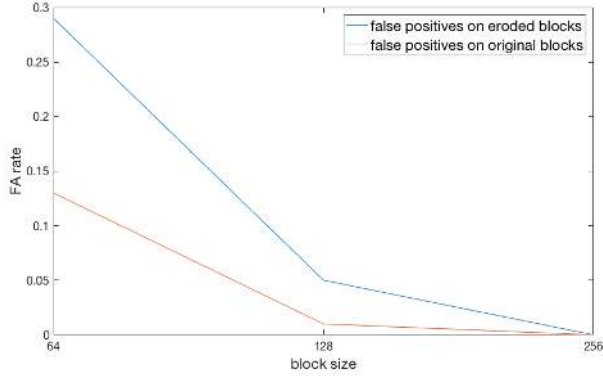


Fig. 7: Dilation detector false alarm (FA) rate for different block sizes.

to the higher probability of having null blocks after erosion. Figures 6 and 7 plot the false positives rates at increasing blocksize, for erosion and dilation detectors, respectively. It is possible to observe that the number of false alarms is a bit lower for the dilation detector, which shows a 0.29% false alarm rate on eroded images at the smallest blocksize, as compared to the almost 2% of the erosion detection on dilated images. The probability of detecting original blocks as filtered are much lower in both cases. More interestingly, the charts show that errors rapidly drop down to zero with the increasing blocksize, leading to a 100% accuracy at 256×256 .

The following set of tests aims at assessing the capability of the detector in discriminating the structuring element B_i used for filtering. In this case, we have a mis-detection if there is a mask not belonging to the decision tree of the applied structuring element that returns a zero value. Since we have already excluded the confusion between filter types (erosion, dilation or null) for a sufficiently large patch, here we focus on the confusion among structuring elements. Experimental tests showed that for both datasets the dilation detector did not produce any mis-detection for every mask and image. This means that the dilation detector was able to classify with 100% accuracy the structuring element used for dilation. On the contrary, the erosion detector generated a few mis-detections in both datasets, mainly associated to larger structuring elements. Tables III and IV report the detailed confusion matrices among

different structuring elements used for erosion on the LRDE and IAM datasets, respectively. Rows are associated to the masks used for filtering, while columns are associated to the masks used for detection. All entries are summed up over the whole dataset, so that each entry (i, j) of the table reports the number of images eroded with B_i and detected as eroded with B_j . The mis-detections are highlighted in red, while the correct detections are highlighted in blue. Please note that the correct detections include the elements implied by Theorem 2 and reported in Table I. As a general consideration, Tables III and IV represent very sparse matrices, meaning that the proposed method allows discriminating the specific mask with a very good average precision. Furthermore, it is to be observed that the values above the diagonal are largely null, except for a couple of exceptions. Consequently, given that our method returns the largest matching element, most of the errors are indeed irrelevant for the final result.

Looking in more detail at red entries in both tables, one can see that the two most critical cases are element B_{26} (diagonal square), and B_{36} (5×5 square), which caused the co-detection of various other non-implied elements. It is to be observed that, although Theorem 2 theoretically establishes the inclusion rules for just a subpart of the elements, non-strict inclusion relationships hold for other element pairs. For instance, the dilation of B_3 by B_4 generates a shape that is very similar, although not exactly equal, to B_{26} , as shown in Figure 8. Accordingly, under erosion with mask B_{26} it cannot be excluded that B_3 and B_4 are sometimes detected, also depending on the characteristics of test image. Since these non-strict implications are in the direction of smaller filters (e.g., the application of B_3 or B_4 will never produce the detection of B_{26}) it is easy to see that selecting the largest detected element guarantees a 99,95% accuracy on the LRDE dataset (2 mis-detections: B_{23} once detected as B_{25} , and B_{32} once detected as B_{34}), and a 100% accuracy on IAM dataset.



Fig. 8: Example of a non-strict inclusion relationship among B_3 , B_4 and B_{26} .

Since there is a substantial symmetry among both filters, it may appear surprising the different behavior of erosion and dilation detector. This could be explained observing that the test images, as usually happens in binary documents, are characterized by thin black structures with horizontal prominence (text, drawings) on white background. This makes the filtered images particularly sensitive to erosion, which sometimes completely erases the objects since the first application, especially for large horizontal structuring elements. This is clearly not true for dilation operators, which does the complementary operation on the background.

ing the result. This causes evident artifacts that can be easily perceived and detected, but obviously destroy most filtering traces, in particular at the object borders. Thus, the possibility of correctly detecting the filter strongly depends on the entity of the rotation or scaling applied to the binary image.

Focusing on the first operation, rotations by $+/- 90$ or 180 degrees, which do not require any interpolation on the image pixels, will be perfectly detected, except for the structuring element returned, which will result rotated by the same amount as the analysed image. This implies 100% accuracy of erosion/dilation detection. Instead, if the filtered image is rotated by an arbitrary angle, morphology is typically used after rotation exactly to regularize and remove aliasing artifacts, thus being correctly detected. In the case of rotation of a filtered image, the traces progressively vanish, but turn out to be persistent to some extent at least for small angles (1-3 degrees). Preliminary analysis performed on a significant subset of images randomly chosen from the two datasets demonstrated that the percentage of non-null pixels in the XOR image generated at the detector for a rotated document is significantly lower if the image was previously filtered, thus allowing a detection. In this case, however, the deterministic properties of the algorithm are lost and a statistical analysis is required, which is out of the scope of the present work. The filtering traces disappear almost completely for larger rotations.

For what concerns rescaling, it is to be mentioned that there is no unique way of performing it on binary images. Typically, after transformation in the grayscale domain, it requires resampling and some filtering or interpolation, followed by further binarization. The above operations, but in particular the filtering, cancel most of the traces even for small rescaling factors. It is to be pointed out that the application of rescaling to binary documents may cause visible artifacts, thus leading to a more realistic scenario in which morphological filtering is applied as a post-processing, in order to conceal the image degradation.

Mirroring is very similar to a 180 degree rotation and is perfectly detectable, except for a corresponding mirroring of the detected structuring element.

A final note on the detector robustness concerns possible malicious manipulations of either the structuring element or the filtered image, with the purpose of faking the detector. Concerning the modification of the mask, as already pointed out in the false alarms analysis, the proposed detector is sensitive to *similar* masks. Accordingly, the application of a structuring element B_{P_i} generated as a small perturbation of a probe element B_i , will leave significant traces of the original mask B_i . As in the above cases, the result will be no more deterministic, in the sense that false negatives will be possible. Nevertheless, a suitable statistical analysis may reveal the filter. On the other side, it is to be pointed out that larger perturbations of the mask (3 or more pixels) will produce significant artifacts on the filtered image, in fact changing the nature of the filter itself and being therefore unusable in practical situations.

The malicious modification of the image is a different situation. It can be modeled as the application of a binary

noise to the filtered image, where pixels at random positions are set to their 1-complement. It is easy to see that each noisy pixel will produce at most a single non-null value in the XOR image at the detector (see Fig. 2), thus leading to a XOR sum equal to or lower than the number of altered pixels. On the other hand, it is to be observed that the impact of such noise on the image is very strong. As a matter of fact, one can measure the PSNR of an attacked binary image I_A of size $N \times M$ as a function of the number of modified pixels P_{mod} , as follows:

$$PSNR(I_A) = 10 \log \left(\frac{N \times M}{P_{mod}} \right). \quad (12)$$

Eq. (12) shows that the modification of even a small percentage of the image points significantly degrades the image. For instance, altering 10000 pixels on a LRDE test image (around 0,1% of the image size) would lead to a PSNR of about 29 dB, with clearly visible artifacts. On the other hand, the sum of the XOR image will remain more than one order of magnitude below the average value returned by a non-filtered image, thus being easily detected with a simple thresholding.

D. Splicing detection

As far as non-geometrical manipulations are concerned, the most important case to consider is the splicing of (part of) two or more different binary images, which is very interesting from an application point of view (e.g., signature substitution or insertion, text manipulation). There are different possible situations to consider. The simplest case is the application of the filter to remove the traces of the insertion of the patch in an image. In this case, the morphological filter is applied after the composition of the two images, thus, it is perfectly detected by the standard method. A more complex situation happens when the filter is applied locally, or when the composition is made by combining original and filtered images, or images processed with different morphological operators. In all these situations, different detection results will be returned for different image subparts according to the applied processing.

Since we have to detect local patches, we need to apply the original algorithm to local areas. To this purpose, we define an analysis window of a given dimension and scan the whole image, applying a classical block-wise analysis. The results output by the detector for each window position form a detection map. If the windows encompasses a null block, the corresponding point in the map is marked as *undetectable*. If the block is not null, the detector is applied: in case of positive result, the relevant point is marked as *possibly filtered*, with the indication of the relevant structuring element; in case of negative result, it is marked as *non-filtered*. At the end of the scan, the resulting detection map is analysed in order to identify significant patches. To this purpose, we search for connected clusters of blocks marked as *possibly filtered* and associated to the same structuring element, and we identify a cluster as a *patch* if it reaches a minimum predefined dimension. All the *possibly filtered* blocks that are not associated to a patch at the end of this process are discarded as noise. In the tests presented in this paper, the image was scanned with non-overlapping square windows of

64×64 pixels, and the minimum patch size was set to 25 connected blocks, corresponding to about 1 square inch at the given resolution of 300 dpi.

In order to create more complex and less visible splicing manipulations, part of the test cases were created by using open and close operators instead of simple erosions and dilations. In fact, open and close do not significantly modify the dimensions of objects, thus being less perceptually relevant. As stated in the introduction, any morphological filtering chain can be simplified in the form of erosions and dilations with composite masks, and the proposed detector is able to reveal the last operation in the chain. In the case of open and close, therefore, the detector will output a dilation or an erosion, respectively. Further analysis would be needed to disambiguate among those operators.

We created a set of 25 test images with different kinds of splicings, by inserting one or multiple patches after erode/dilate/open/close operations with different structuring elements. The patch size was set to 320×640 pixels, corresponding to 1×2 inches at the given resolution of 300 dpi. According to the above described procedure, the detection algorithm produced for each test image a detection map, in which each block was marked as either undetectable, non-filtered, or filtered. Each cluster of connected filtered blocks recognized as a patch was then marked with a unique label and annotated with the relevant structuring element.

The first important result is that the algorithm was able to correctly reveal the presence of the filtered patch for all the splicings and for any applied filter (operator and mask). Furthermore, no false positives were returned for any of the test images, thus achieving 100% accuracy in splicing detection. As far as the localization of the manipulated patch is concerned, the adopted block analysis cannot guarantee a full resolution at the patch boundaries, due to the possible misalignment of patch and blocks. Accordingly, 94% of the blocks overlapped with the patch were correctly detected, with all the errors located at the patch border. As already discussed in Section III.B, also in this case the statistics of dilation (or open) and erosion (or close) detectors turned out to be different, with the first performing slightly better (95% vs. 93%). To achieve higher localization precision, possible solutions may include performing a hierarchical refinement at block boundaries, or using an overlapped sliding window to adjust the border of the patch.

Figures 9 and 10 show two splicing test images (original on the left, spliced in the center) as well as the relevant detection maps (on the right). Each image contains two patches filtered with different masks. In the first example the patches are filtered with an open operator with masks B_3 and B_4 , respectively. The second example refers to a close filtering with masks B_3 and B_{15} , respectively. Consequently, the reported maps were output by the dilation detector for the former and by the erosion detector for the latter. The reciprocal detectors did not produce any output. In both cases it is possible to observe that the detector identifies all the manipulated patches. In the case of open, the two spliced filtered patches are perfectly detected. The missing blocks in the detected rectangles are indeed null blocks, correctly

classified as undetectable. In the case of close, some errors are instead present at the block boundary.

IV. CONCLUSIONS

In this paper we have presented a new approach to detect the application of morphological filters to binary images. The importance of binary document forensics has been highlighted, underlining the complete lack of techniques able to address the detection of this family of filters, largely used in document image manipulation. The basic theory of morphological filters has been reviewed, focusing on fundamental properties of the basic operators, erosion and dilation, and deriving some theoretical consequences that could be exploited to find the traces of the filters. Then, a forensic algorithm has been proposed to perform the detection of a morphological operator, as well as the relevant structuring elements used. Finally, the results of an extensive testing are presented, showing that the proposed approach is highly effective, as well as robust to the typical manipulations of binary image documents, including splicing. Future directions of this work could lead to the detection of multiple filterings (e.g., open and close operators and filter chains). Furthermore, the possibility of using the proposed technique as a basis for a statistical analysis of filtering traces may open the possibility of dealing with more severe attacks such as large rotations and scaling, as well as extending the technique to graylevel image morphology, where compression may become a critical issue.

REFERENCES

- [1] P.-J. Chiang, N. Khanna, A. K. Mikkilineni, M. V. Ortiz Segovia, J. P. Allebach, G. T. C. Chiu, and E. J. Delp, "Printer and scanner forensics: models and methods," in *Intelligent Multimedia Analysis for Security Applications*. Studies in Computational Intelligence, vol. 282, 2010, pp. 145–187.
- [2] E. J. Delp, N. Memon, and M. Wu, Eds., *Special Issue on Forensics analysis of digital evidence*, *IEEE Signal Processing Mag.*, vol. 26, no. 2, 2009.
- [3] A. Khanna, A. K. Mikkilineni, A. F. Martone, G. N. Ali, G. T. C. Chiu, J. P. Allebach, and E. J. Delp, "A survey on forensics characterization methods for physical devices," *Digital Investigations*, vol. 3, pp. S17–S28, 2006.
- [4] H. Farid, *Photo Forensics*. MIT Press, 2016.
- [5] K. Bahrami, A. Kot, L. Li, and H. Li, "Blurred image splicing localization by exposing blur type inconsistency," *IEEE Trans. on Information Forensics and Security*, vol. 10, no. 5, pp. 999–1009, 2015.
- [6] G. Cao, Y. Zhao, R. Ni, and X. Li, "Contrast enhancement-based forensics in digital images," *IEEE Trans. on Information Forensics and Security*, vol. 9, no. 3, pp. 515–525, 2014.
- [7] G. Cao, Y. Zhao, R. Ni, and A. Kot, "Unsharp masking sharpening detection via overshoot artifacts analysis," *IEEE Signal Processing Letters*, vol. 18, no. 10, pp. 603–607, 2011.
- [8] M. Kirchner and J. Fridrich, "On detection of median filtering in digital images," in *SPIE Media Forensics and Security II*, vol. 7541, (2010), pp. 101–112.
- [9] G. Cao, Y. Zhao, R. Ni, L. Yu, and H. Tian, "Forensic detection of median filtering in digital images," in *IEEE Int. Conf. on Multimedia and Expo, ICME2010*, (2010), pp. 89–94.
- [10] H.-D. Yuan, "Blind forensics of median filtering in digital images," *IEEE Trans. on Information Forensics and Security*, vol. 6, no. 4, pp. 1335–1345, 2011.
- [11] X. Kang, M. C. Stamm, A. Peng, and K. J. R. Liu, "Robust median filtering forensics using an autoregressive model," *IEEE Trans. on Information Forensics and Security*, vol. 8, no. 9, pp. 1456–1468, 2013.

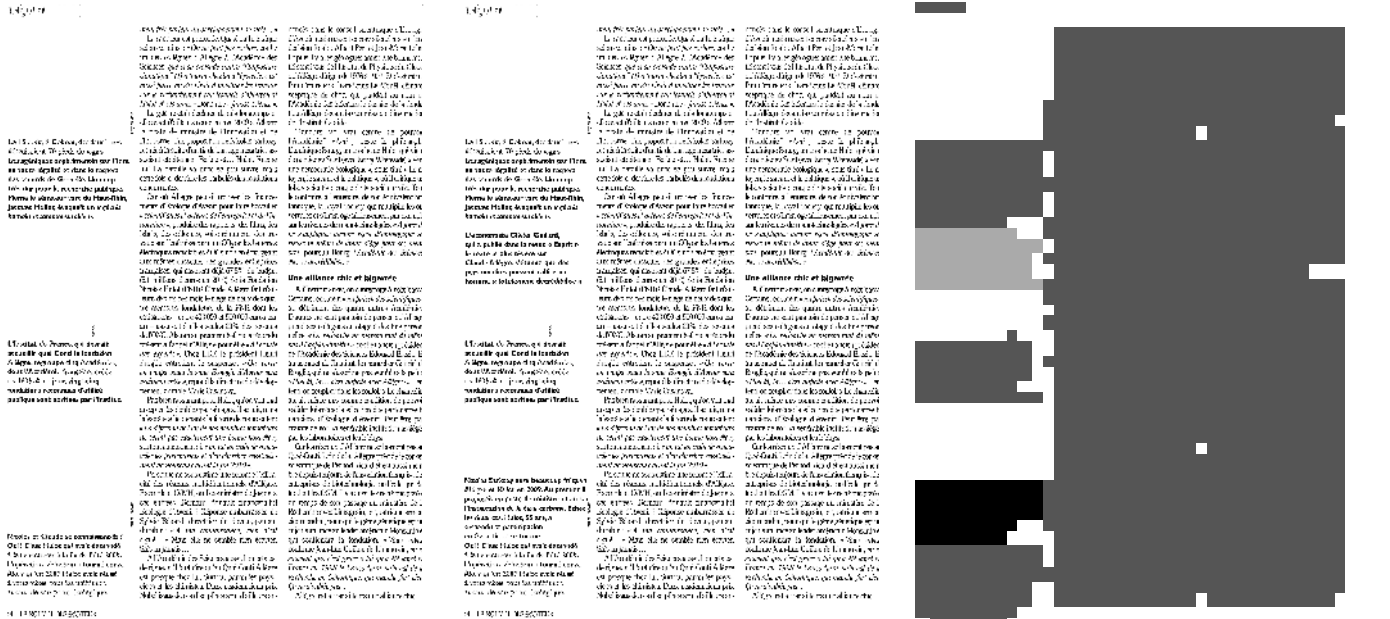


Fig. 9: Example of splicing detection: (left) LRDE test image; (center) spliced image with two patches of size 320×640 pixels after open operation with B_3 and B_4 , respectively; (right) dilation detection map. Undetectable blocks are marked in white, non-filtered blocks in dark gray, filtered blocks with element B_3 in light gray, and filtered blocks with element B_4 in black.

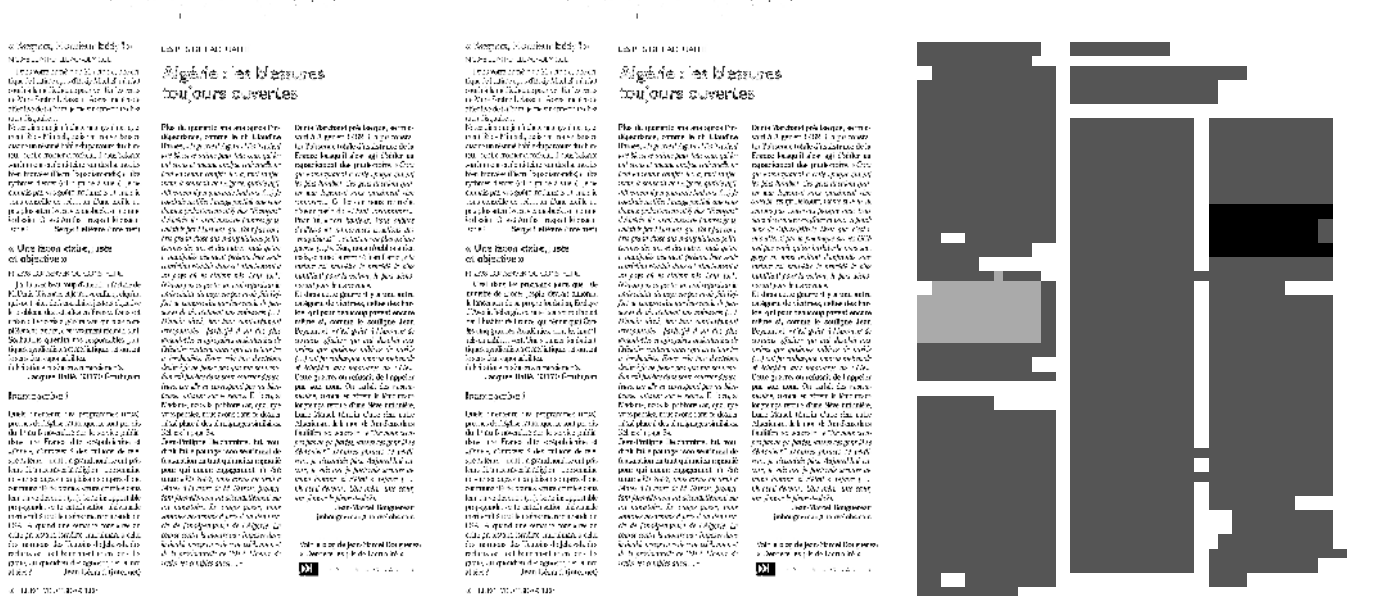


Fig. 10: Example of splicing detection: (left) LRDE test image; (center) spliced image with two patches of size 320×640 pixels after close operation with B_3 and B_{15} , respectively; (right) erosion detection map. Undetectable blocks are marked in white, non-filtered blocks in dark gray, filtered blocks with element B_3 in black, and filtered blocks with element B_{15} in light gray.

- [12] C. Chen, J. Ni, and J. Huang, "Blind detection of median filtering in digital images: A difference domain based approach," *IEEE Trans. on Image Processing*, vol. 22, no. 12, pp. 4699–4710, 2013.
- [13] Y. Zhang, S. Li, S. Wang, and Y. Q. Shi, "Revealing the traces of median filtering using high-order local ternary patterns," *IEEE Signal Processing Letters*, vol. 21, no. 3, pp. 275–280, 2014.
- [14] J. Chen, X. Kang, Y. Liu, and Z. J. Wang, "Median filtering forensics based on convolutional neural networks," *IEEE Signal Processing Letters*, vol. 22, no. 11, pp. 1849–1853, 2015.
- [15] L. Zhang, W. Duan, and H. Guo, "Fuzzy operation forensics research based on mathematical morphology," *Lecture Notes in Electrical Engineering*, vol. 154, pp. 1457–1462, 2012.
- [16] W. Luo, J. Huang, and G. Qiu, "A novel method for block size forensics based on morphological operations," *Digital Watermarking, Lecture Notes in Computer Science*, vol. 5450, pp. 229–239, 2009.
- [17] C. Pasquini, G. Boato, N. Anaijlan, and F. G. B. De Natale, "A deterministic approach to detect median filtering in 1D data," *IEEE Trans. on Information Forensics and Security*, vol. 11, no. 7, pp. 1425–1437, 2016.
- [18] A. Haas, G. Matheron, and J. Serra, "Morphologie mathématique et granulométries en place," *Annales de Mines*, vol. 11, no. 12, pp. 736–753, 1967.
- [19] F. Y. Shih, *Image processing an Mathematical Morphology: Fundamentals and Applications*. CRC Press, 2009.
- [20] Y. Nakagawa and A. Rosenfeld, "A note on the use of local min and max operations in digital picture processing," *IEEE Trans. on Systems, Man and Cybernetics*, vol. 8, no. 8, pp. 632–635, 1978.
- [21] V. Serra, "An overview of morphological filtering," *Circuits, Systems and Signal Processing*, vol. 11, no. 1, pp. 47–108, 1992.
- [22] G. Lazzera and T. Geraud, "Efficient multiscale Sauvola's binarization," *Int. Journal on Document Analysis and Recognition*, vol. 17, no. 2, pp. 105–123, 2014.
- [23] U. Marti and H. Bunke, "The IAM-database: An english sentence database for off-line handwriting recognition," *Int. Journal on Document Analysis and Recognition*, vol. 5, pp. 39–46, 2002.



Francesco De Natale (M.Sc. '90, Ph.D. '94) is a Professor of Telecommunications at the University of Trento, Italy, where he leads the Multimedia Lab (mmlab.disi.unitn.it). His research interests are focused on multimedia communications, with particular attention to multimedia signal processing, analysis, and retrieval. He was Program Co-Chair of the IEEE Intl. Conf. on Image Processing (ICIP-2005) and General Chair of the ACM Intl. Conf. on Multimedia Retrieval (ICMR-2011). He has been Associate Editor of the IEEE Trans on Multimedia

and of the IEEE Trans. on Circuits and Systems for Video Technologies, as well as a member of the IEEE Signal Proc. Society Technical Committee on Multimedia Signal Processing (MMSP), chairing the Technical Directions Subcommittee. Currently, he is member of the Board of Directors of the Italian Consortium for Telecommunications (CNIT), and member of the management of the Italian Group of Telecommunications and Information Technologies (GTTI). He is also co-founder of the university startup Xtensa, a company that develops tools in the area of user interaction and ambient intelligence. He published more than 150 papers in international journals and conferences, mostly in the area of multimedia signal processing and communications, and has been scientific coordinator of many large-scale research and development projects, both at the national and international level. Prof. De Natale was appointed evaluator for several international bodies, including the European Commission and the NSFs of US and Ireland. Prof. De Natale is a Senior Member of IEEE and a member of ACM and GIRPR.



Giulia Boato is Assistant Professor at the Department of Information Engineering and Computer Science (DISI) of the University of Trento (Italy) and professor of the courses Digital signal processing, Multimedia data security and Probability and statistics within the M.Sc. and the B.Sc. Degree in Telecommunications Engineering. From 2008 to 2011 she has been coordinating the Multimedia Signal Processing and Understanding Lab. She was in the Project staff of many projects (FP7 FET-IP LIVINGKNOWLEDGE, FP7 IP GLOCAL, FP7 CA ETERNALS, ICT-COST 3D-ConTourNet). She was co-chair of the International Workshop Living Web: making diversity a true asset (Washington DC, October 2009) within the International Semantic Web Conference 2009 and of the workshop on Event-based Media Integration and Processing co-located with ACM Multimedia conference 2013. She is reviewer for many international journals, e.g., IEEE Transactions on Information Forensics and Security, IEEE Transactions on Signal Processing, IEEE Transactions on Multimedia, IEEE Transactions on Image Processing, IEEE Transactions on Circuits and Systems for Video Technology. She is in the editorial board of the Elsevier Image and Vision Computing Special Issue on Event-based Media Processing and Analysis. Her research interests are focused on image and signal processing, with particular attention to multimedia data protection, data hiding and digital forensics, but also intelligent multidimensional data management and analysis. She is author of 85 papers in international conferences and journals. She is member of the IEEE Multimedia Signal Processing Technical Committee (MMSP TC).