

Scanned documents forensics: detecting inserted characters through noise and chromatic artifacts

Marina Gardella^{1*}, Julieta Umpierrez¹, Antoine Tadros^{1,2}, Seginus Mowlavi¹, Natalia Bottaioli¹, Diego Belzarena^{1,3}, Gabriele Facciolo¹, Roy Y. He⁴, Jean-Michel Morel⁴, Rafael Grompone von Gioi¹

¹ *Université Paris-Saclay, ENS Paris-Saclay, CNRS, Centre Borelli, 91190, Gif-sur-Yvette, France*

² *Déterminant, France*

³ *Instituto de Ingeniería Eléctrica, Facultad de Ingeniería, Universidad de la República, Uruguay*

⁴ *City University of Hong Kong, Hong Kong*

*marina.gardella@ens-paris-saclay.fr

Abstract

Document forgery detection plays a crucial role in safeguarding the integrity of various sectors, including accounting, insurance, finance, law enforcement, and national security. The ability to distinguish between authentic and counterfeit documents is key to ensure trust in transactions, maintain regulatory compliance, and prevent fraudulent activities. Despite its importance, the field of document forgery detection remains underdeveloped, especially compared to the advances made in image forgery detection. In this work, we present the first benchmarking results of state-of-the-art image forgery detection methods applied to forged documents, and we demonstrate that these methods underperform in this context. To address this gap, we introduce two novel approaches specifically designed for document forgery detection. These approaches analyze anomalies in noise distribution and chromatic artifacts in scanned documents to identify inserted characters. The source code is available in <https://github.com/julietaumpierrez/Scanned-documents-forensics>

1. Introduction

Assessing the integrity of a document is essential for ensuring trust in transactions, maintaining regulatory compliance, and preventing fraudulent activities. To address document integrity, one possible approach is to identify the device responsible for generating the digitized or scanned image [10]. Forensic experts can verify the authenticity of a document and possibly identify its owner by associating an image with the source device. This process helps establish a chain of custody and authenticity. Another approach is to verify if the content of the document has been altered. In this case, the goal is to detect potential falsification, ascertain the type of alteration, and localize the tampered regions.




Figure 1. (a) Genuine and (b) inserted characters zoomed-in from a forged document. The inserted character lacks the sensor noise and chromatic artifacts present in the genuine one.

Despite its significance, the field of document forgery detection remains underdeveloped, especially when compared to the progress in image forgery detection. While scanned documents can be treated as images, the distinct processing pipeline of scanners poses challenges when directly applying traditional image analysis techniques. Forensic methods tailored to documents should address the complexities found across different document types.

This work presents two primary contributions. First, to the best of our knowledge, we conduct the first comprehensive benchmarking study of state-of-the-art image forgery detection methods on a dataset of forged scanned documents. Our evaluation reveals that these methods consistently underperform in this context, highlighting a critical gap in their generalizability and the need for domain-specific solutions in document forensics.

Second, we propose two novel approaches explicitly tailored for detecting forgeries in scanned documents, both

aimed at identifying inserted characters through the analysis of inconsistencies in scanner-related traces. The first approach focuses on the analysis of local noise distributions. Genuine characters exhibit sensor noise introduced during the acquisition process. In contrast, digitally inserted ones often lack this noise, enabling their detection through statistical noise modeling. The second approach leverages chromatic aberration—an optical distortion typically introduced by the scanner’s lens system. While genuine characters are subject to this artifact, inserted ones are not, resulting in detectable chromatic inconsistencies. These visual anomalies, illustrated in Fig. 1, form the foundation of our detection pipeline. Both proposed methods are designed to maintain a low false positive rate while effectively distinguishing between authentic and manipulated content.

The rest of the document is organized as follows: Sec. 2 presents the related work, both for image and document forgery detection. Sec. 3 presents both proposed approaches. Sec. 4 presents the benchmarking details and evaluation results. Sec. 5 presents the proposed follow-up work and finally Sec. 6 presents this work’s conclusion.

2. Related works

The task of detecting forged content in digital documents lies at the intersection of image forensics and document analysis. While extensive research has been devoted to general-purpose image forgery detection, far less attention has been paid to methods tailored specifically for documents, particularly scanned ones. In this section, we review prior work in image forensics, highlight its limitations in document scenarios, and summarize efforts specifically targeting document forgery detection.

2.1. Image forgery detection

There is a large literature on image forensics, starting from the seminal work of Farid [20]. Classical methods focus on identifying local anomalies in the traces left by the digital camera processing chain. These local anomalies suggest the presence of a forgery. Such forgery detection methods can be classified by the specific traces they target, including noise-based approaches [21, 32, 33], CFA-based methods [3, 11, 45] and compression-based techniques [29, 38].

Generic tools based on neural networks were recently proposed to detect complex inconsistencies. Splicebuster [16] detects inconsistencies in a residual and more recently TruFor [23] utilizes a fusion architecture based on transformers, combining RGB images and a learned noise-sensitive fingerprint called Noiseprint++ [13]. CAT-Net [26] is an end-to-end convolutional neural network fusing an RGB stream and a discrete cosine transformation stream. PSCC-Net [30] leverages both spatial and channel-wise correlation features at different scales. FOCAL [50] clusters features acquired from contrastive learning. EXIF as Language [51]

proposes a multi-modal embedding of image patches and EXIF metadata, using clustering in this space for forgery detection.

As illustrated in Fig. 2, applying traditional image forensic methods to digitized documents poses several challenges. Indeed, high-quality scans are typically saved in lossless formats [24], making compression-based methods ineffective. Noise-based methods are susceptible to false detections when dealing with textures. Furthermore, saturated backgrounds can also cause false detections due to noise clipping in those areas. CFA-based methods are designed to detect anomalies in the 2×2 color filter arrays found in digital cameras. While sensors in scanners also have color filter arrays, they are trilinear instead [4]. Some generic tools such as PSCC-Net [30] or TruFor [23] seem to adapt better to documents. However, they highlight more and larger areas than those being altered. Since these examples illustrate potential limitations, a more systematic evaluation is necessary. In Sec. 4, we present a comprehensive benchmark to assess the performance of these methods on a dataset of forged documents.

2.2. Document forgery detection

One line of research in document forgery detection relies on advanced imaging techniques such as hyperspectral imaging [22, 44, 46], laser-induced breakdown spectroscopy [27], Terahertz imaging [47], X-Ray Fluorescence [35, 41] and Raman Spectroscopy [7, 12]. These techniques analyze the composition and properties of the materials, such as ink or paper. Although their performance is good, they require access to the physical document and the use of sophisticated equipment making them hard to recommend for large-scale forgery detection.

As the volume of digital documents continues to grow along with the accessibility of image manipulation tools, digitally altered documents have become increasingly common. In response to this trend, the *Find it! Fraud Detection Contest* was launched during the 2018 ICPR Contest Session [1]. The competition challenged participants to detect and localize forgeries in digital documents. Five teams participated in the detection task, while two addressed the more demanding localization task.

The competing teams employed diverse strategies. Text-based approaches assessed the internal consistency of document contents, such as verifying calculations and date logic. Image-based methods, on the other hand, leveraged techniques such as copy-move forgery detection [15], Noiseprint [13], and StegoFeatures [14]. These methods, adapted from the image forgery detection field, suffer from the drawbacks pointed out in the previous section.

Beyond these, a number of works have proposed detecting forgeries by exploiting structural and visual inconsistencies specific of documents. Bertrand et al. [5] propose a method

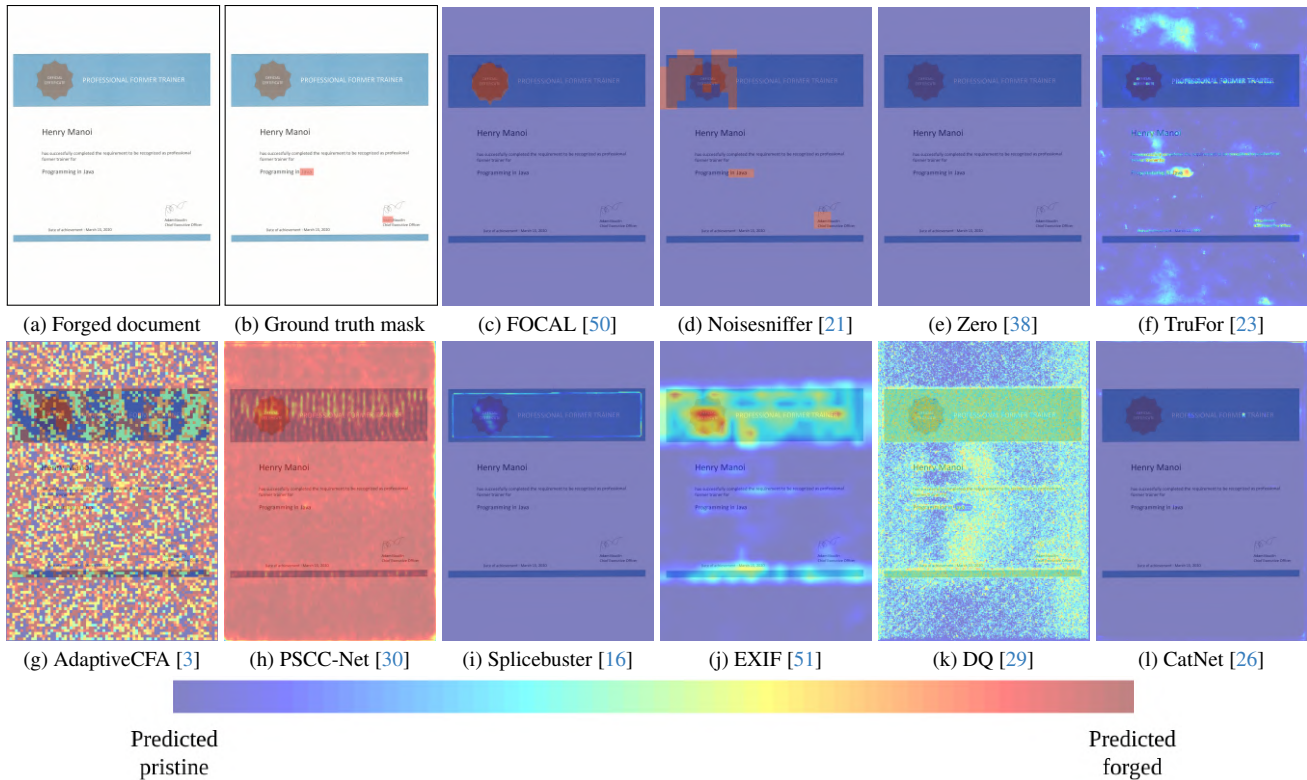


Figure 2. Example of outputs with common image forgery detection methods for image `s9_52_c` of the SUPATLANTIQUE [43] dataset with corresponding ground truth. Note that most methods make incorrect predictions.

to detect outlier characters based on features such as size, inertia axis and alignment. These features are then fed to a support vector machine trained for classification. This idea was extended in [6] to identify font inconsistencies via a conditional random field framework, based on the computation of conditional probabilities. More recent work by Joren et al. [25] leverages OCR-extracted features to obtain a graph-based representation that is later classified using a random forest. Texture-based cues have also been explored. For instance, Cruz et al. [17] apply local binary patterns to identify background inconsistencies in patches. This idea is also present in the work by Centeno et al. [8] where they propose several texture descriptors to leverage texture inconsistencies in banknotes and identity cards. Tornés et al. [48] propose using knowledge-graph embeddings that directly search for semantic discrepancies in the text and later in [49] they use natural language processing techniques to tackle the problem in a similar way. Nandanwar et al. [36] propose using the discrete cosine transform in order to obtain features that are later used as an input to a convolutional neural network that classifies the document as forged or pristine.

Recently, deep neural network-based tools have been proposed to detect inconsistencies in documents. There have also been some works such as [42] that propose deep learning based methods to detect fake text in images taken with a com-

mon camera. Dong et al. [19] propose an encoder-decoder architecture with multi-scale attention modules. Chen et al. [9] introduce a neural network that combines a visual enhancement module with a wavelet-like frequency enhancement module to perform document forgery detection. Luo et al. [31] use multi-modal information fusion and contrastive learning. Li et al. [28] propose using spatial-frequency domain features with a HRNet-inspired architecture to detect different kinds of forgeries in documents. Nguyen et al. [37] propose TALIU, a lightweight segmentation-based model.

Collectively, these works introduce novel and promising approaches to document forgery detection. While earlier methods rely on hand-crafted features combined with traditional classifiers such as support vector machines, more recent approaches follow a deep learning paradigm that learns representations directly from data. To varying degrees, all of these methods are data-driven, requiring annotated examples to train models or tune parameters. However, most of them have neither released the code nor the training datasets which limits their reproducibility and makes them unsuitable for benchmarking or broader public use.

3. Proposed approaches

We propose two methods to detect inserted characters in scanned documents. Each method exploits a different sensor

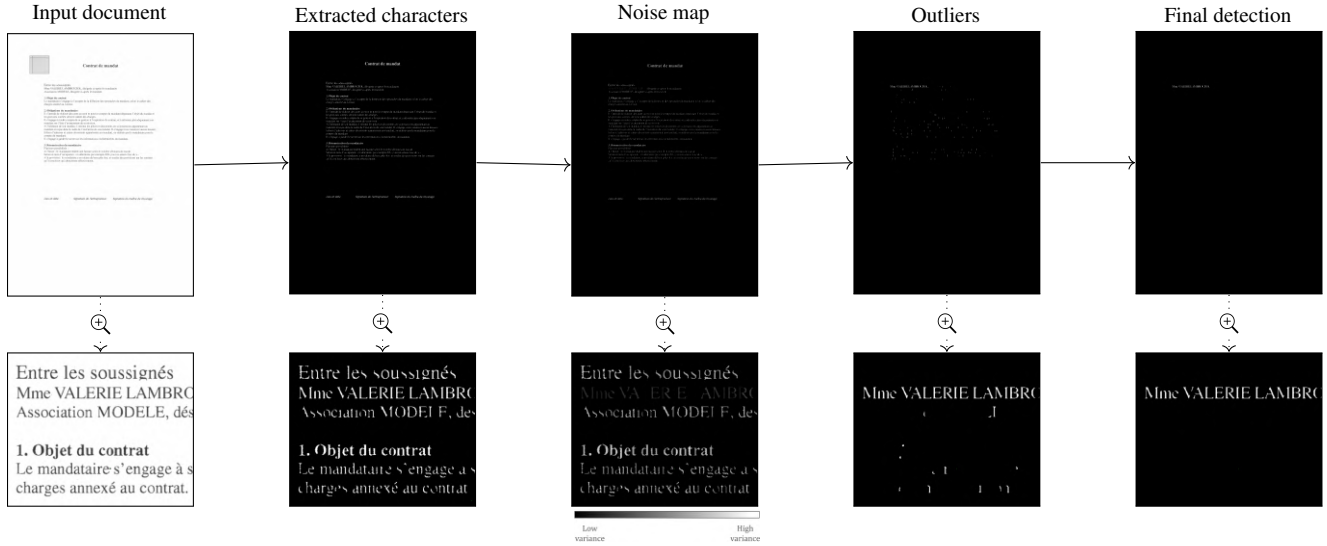


Figure 3. Pipeline of the document forgery detection method based on noise analysis. The top row shows the main processing steps: (1) the input document image, (2) the extracted characters, (3) a map of the local noise level for each character, (4) characters identified as noise outliers, and (5) the final detection result highlighting suspicious words. The bottom row presents zoomed-in views of each step for a selected region. Dotted arrows connect each full-resolution image with its corresponding zoomed view. A noise colormap is shown below for reference.

artifact related to the scanning process. The first approach analyzes the noise present in the image of the characters' pixels. Inserted characters lack the typical intensity fluctuations in the pixels that constitute them and which can be seen in genuine ones. The second approach leverages chromatic artifacts. Genuine characters show color fringing, unlike inserted ones, as illustrated in Fig. 1.

3.1. Noise based forgery detection

Let $\mathcal{C} = \{c_1, \dots, c_N\}$ be a list of N characters extracted from a document. Each character $c_i \in \mathcal{C}$ is represented by a collection of pixels. These pixels can be easily separated from the background by thresholding. Let s_i denote the standard deviation of character c_i , defined as the standard deviation of its pixels' intensities, with $i = 1, 2, \dots, N$.

Given a list of characters and their corresponding standard deviations, we build a prediction interval for s_1, \dots, s_N . To do so, we compute the mean μ_s and the standard deviation σ_s of the samples. Given a confidence level $0 \leq \alpha \leq 1$, we define the prediction interval as

$$\mathcal{I}_\alpha = [\mu_s - z_\alpha \sigma_s; \mu_s + z_\alpha \sigma_s], \quad (1)$$

where

$$\mu_s = \frac{\sum_{i=1}^N s_i}{N}, \quad (2)$$

$$\sigma_s^2 = \frac{\sum_{i=1}^N (s_i - \mu_s)^2}{N - 1}, \quad (3)$$

and z_α is the critical value such that $P(Z > z_\alpha) = \alpha$ for $Z \sim \mathcal{N}(0, 1)$.

Here, we are assuming that s_1, \dots, s_N follow a Gaussian distribution $\mathcal{N}(\mu_s, \sigma_s)$. Assuming that the noise is homoscedastic and Gaussian, the standard deviation at each pixel follows a χ -distribution with one degree of freedom χ_1 . When computing the variance of a character, we add up several independent χ^2 -distributions with one degree of freedom, resulting again in a χ^2 -distribution where the degrees of freedom are given by the number of pixels in the character. By the central limit theorem, since the number of pixels in a character is usually large, we can approximate this χ -distribution by a Gaussian.

To detect potential outliers, we define the subset $\mathcal{L}_\alpha \subset \mathcal{C}$ of those characters whose standard deviation is smaller than the lower bound of \mathcal{I}_α . Namely,

$$\mathcal{L}_\alpha = \{c_i \in \mathcal{C} : s_i < \mu_s - z_\alpha \sigma_s\}. \quad (4)$$

Note that \mathcal{L}_α can contain inserted and genuine characters.

We filter out genuine characters using the *a contrario* methodology [18] based on the non-accidentalness principle [2]. Let *words* be the elements of a given partition of \mathcal{C} and let the length of a word be its cardinal. Given a word W of length n , define

$$K_\alpha(W) = \sum_{c \in W} X_\alpha(c) \text{ with } X_\alpha(c) = \begin{cases} 1 & \text{if } c \in \mathcal{L}_\alpha \\ 0 & \text{if } c \notin \mathcal{L}_\alpha \end{cases}. \quad (5)$$

Under the null hypothesis (H_0) that all variables in $\{X_\alpha(c) \mid c \in W\}$ are independent and follow a Bernoulli distribution with parameter α , $K_\alpha(W)$ follows a binomial distribution $\text{Bin}(n, \alpha)$. The number of false alarms (NFA) [18] is defined by

$$\text{NFA}(W) = N_T \mathcal{P}(Z \geq K_\alpha(W)), \quad (6)$$

where $Z \sim \text{Bin}(n, \alpha)$ and the number of tests $N_T \in \mathbb{N}$ is the number of words in the document. The NFA is an upper bound on the expectation of abnormal occurrences under H_0 . For an *a priori* estimate of the mean number $\varepsilon > 0$ of false detections under H_0 , a word W is said to be ε -meaningful if $\text{NFA}(W) < \varepsilon$ holds. Once ε is fixed, a word W is detected if it is ε -meaningful. This means that the expected number of words that are ε -meaningful under H_0 is smaller than ε . This pipeline is summarized in Fig. 3.

3.2. Forgery detection based on chromatic artifacts

Our second approach leverages local chromatic fluctuations instead of image noise. These chromatic fluctuations are shown in Fig. 4. We first preprocess the document to highlight its chromatic artifacts. Let $I = (R, G, B)$ be an image with three color channels, red (R), green (G) and blue (B). The proposed preprocessing first subtracts the gray-mean $m_p = (R_p + G_p + B_p)/3$ from each element of the triplet $p = (R_p, G_p, B_p)$ for each pixel p . After this preprocessing, the pixel's values are $\hat{p} = (R_p - m_p, G_p - m_p, B_p - m_p)$. If a pixel p was originally grayscale ($R_p = G_p = B_p$), then $\hat{p} = (0, 0, 0)$. If the pixel has chromatic artifacts, \hat{p} will show variations around the mean of the three channels.

We now apply the detection strategy in Sec. 3.1 by replacing the per-character standard deviations s_i in (4) with the one computed on the preprocessed pixels \hat{p} .

3.3. Implementation details

Binarization is done using a multi-class extension of Otsu's method [39], which provides a fast and unsupervised thresholding baseline suitable for general-purpose document processing. For the noise-based approach, we use a 4-class Otsu segmentation. The goal is to isolate the core of each character while avoiding the transitional regions that tend to increase the estimated intensity variance, which could interfere with our statistical modeling. For the chromatic-based approach, we apply a 3-class Otsu method. While we still aim to reduce the impact of transitional regions, in this case, we deliberately preserve the chromatic artifacts surrounding characters. In both cases, class selection is designed to avoid overestimating character variance while preserving features that are critical for forgery detection.

To extract characters from a given document, we perform connected component analysis. This procedure is done channel-wise and the final list of characters is obtained as the intersection of the channel-wise ones. This helps reduce false

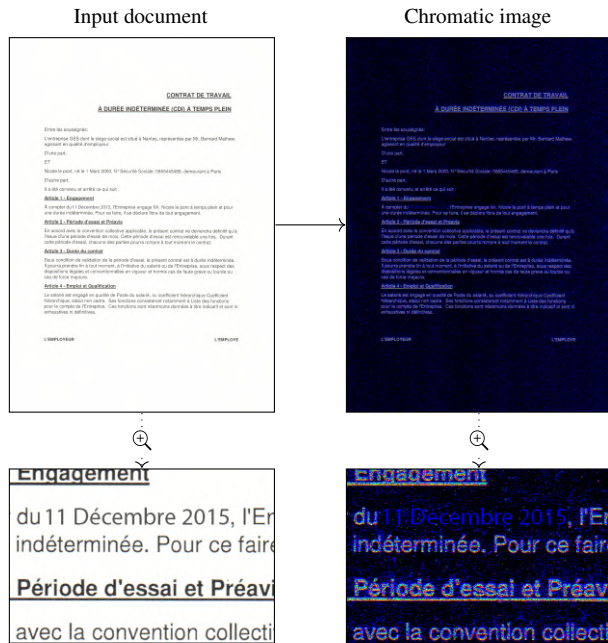


Figure 4. Input document (left) and chromatic image (right). The preprocessing step removes the grayscale component from each pixel, making subtle color variations—such as chromatic fringes—more visible. Inserted characters, which typically lack these artifacts, appear noticeably flatter in the chromatic image.

positives due to color-specific noise or artifacts. Characters are usually consistent across channels, so the intersection acts as a noise filter. This list is further refined discarding small blobs of less than 30 pixels. Small connected components are typically noise (e.g., punctuation artifacts, dust, compression artifacts). A minimum size threshold removes such noise without discarding legitimate characters.

To build the prediction interval, we set $\alpha = 0.1$, a common default in statistical hypothesis testing. As for the NFA, we fix the NFA threshold ε to 0.01, which means that we expect at most 1 false detection in every 100 trials under the null hypothesis. Words are extracted using the EAST text detector [52]. EAST offers a good balance between speed and accuracy, making it suitable for large-scale document analysis. The modular framework supports integration of newer or specialized text detectors.

4. Experiments

Evaluated methods. Besides evaluating both of the proposed methods we also include what, to the best of our knowledge, is the first benchmarking of image forgery detection methods in a forged document dataset. The chosen methods are Noisesniffer [21] which detects anomalies at noise level, Zero [38] which detects anomalies in the JPEG grid, FOCAL [50], EXIF [51], CAT-Net [26], Adaptive [3], PSCC-Net [30], TruFor [23], which are all deep

learning based methods, Splicebuster [16] which uses expectation maximization to analyze a high frequency residual and DQ [29] which finds double compression traces. With the exception of Noisesniffer, Zero and FOCAL, the rest of the methods have a heatmap for output. In that case we binarized them using a threshold of 0.5. These methods were evaluated on the PhotoHolmes library [40].

Dataset. All methods were evaluated on the SUPATLANTIQUE dataset [43]. This dataset was built using 11 flatbed scanners of different brands and models. Three common forgeries were applied to these scanned documents: copy-move, splicing, and retouching. We only used the retouching dataset as it is the one that contains our target forgery: inserted characters. This dataset contains 34 images of size 2481×3471 . One could think that the size of this dataset is small. However, it is important to keep in mind that our analysis is done at the word level and the dataset contains 6229 words, which yields a sufficient number of samples for a fair evaluation.

Metrics. We evaluated the methods in terms of precision, recall, accuracy, and $F1$ -score, defined as

$$\text{precision} = \frac{TP}{TP + FP}, \quad (7)$$

$$\text{recall} = \frac{TP}{TP + FN}, \quad (8)$$

$$\text{accuracy} = \frac{TP + TN}{TP + TN + FP + FN}, \quad (9)$$

$$F1\text{-score} = \frac{2 \times TP}{2 \times TP + FP + FN}. \quad (10)$$

Since our level of analysis is words, metrics are computed at the word level for each document. A word is forged (positive class) if any of its characters is. Conversely, a word is genuine (negative class) if none of its characters are forged. Once the scores are computed on each document, we average all the results in order to obtain a single value that reflects the performance over the whole dataset.

Results. The results obtained for both of our proposed approaches, along with the benchmarked image forgery detection methods, are presented in Tab. 1.

We observe that both of the proposed methods achieve a high precision. This is due to the fact that both methods deliver very few false detections, if any. In particular, the noise method delivers no false positives while the chromatic one delivers two. Accuracy is also high for both methods. However, due to the imbalanced nature of the dataset (i.e. there are more genuine words than forged words), this is explained mainly by the fact that both methods are able to correctly predict negative samples. In terms of recall and $F1$ -score, we note that the chromatic method almost doubles the scores obtained by the noise method. Although

Method	$F1$ -Score	Precision	Accuracy	Recall
Noise (Sec. 3.1)	0.3485	1.0000	0.9544	0.2988
Chromatic (Sec. 3.2)	0.6286	0.9608	0.9584	0.5684
Noisesniffer [21]	0.1511	0.3100	0.8717	0.3343
Zero [38]	0.0000	1.0000	0.9462	0.0000
FOCAL [50]	0.0599	0.1255	0.4617	0.5838
EXIF [51]	0.0000	0.9118	0.9405	0.0000
CAT-Net [26]	0.0000	0.8235	0.9418	0.0000
Adaptive [3]	0.0744	0.0530	0.0623	0.9893
PSCC-Net [30]	0.0313	0.3963	0.5514	0.4344
TruFor [23]	0.2828	0.4850	0.8921	0.3676
Splicebuster [16]	0.0911	0.6559	0.9401	0.1225
DQ [29]	0.0000	1.0000	0.9462	0.0000

Table 1. $F1$ -score, precision, accuracy, and recall of the proposed approaches on the SUPATLANTIQUE [43] dataset considering only forged images from the retouching dataset.

the chromatic method delivers two false detections, its ability to predict positive samples is better.

Regarding the classical forgery detection methods, none of them deliver good results in this new domain. Both of our proposed methods give better results.

Zero, EXIF, CAT-Net and DQ all share similar results: high precision and accuracy with low recall. This is explained by the fact that all of them produce mostly null masks, with Zero and DQ systematically producing masks that predict no forgeries at all in each image. For the likes of Zero, DQ and CAT-Net this result is expected given that the three of them are designed to find inconsistencies delivered by JPEG artifacts. In this case, the forged documents were not compressed, so no inconsistencies of this kind can be found. In the case of EXIF, even though scanned documents do contain EXIF information, the results are bad because of the poor generalization of this method to this type of images that was not included in their training set.

On the other hand, PSCC-Net, FOCAL, Splicebuster and Noisesniffer perform similarly to each other. Their results are explained by how the outputs are not all zeros such as in the previous group, but the predictions they give are not correctly localized. Given that both PSCC-Net and FOCAL are deep learning-based methods, this is explained by their poor generalization to this new domain. In the case of Splicebuster, this method obtains a residual which, because of the saturated areas present in documents, is not suited for this case. In Noisesniffer, performance is hindered by saturated regions in documents, as the method discards them.

On the other side of the spectrum of results, we have Adaptive. Having a high recall but low precision means that the method is systematically giving predictions that flag a

huge part of the image as forged. This bad performance is expected, as this method tries to predict inconsistencies in the Bayer array while scanners use a linear array instead.

Finally, even if TruFor is considered the state-of-the-art in image forgery detection, the results are not as good as the ones reported for the image forgery detection task. This can be explained by the Noiseprint++ extraction being affected by saturated areas present in documents. Despite some good detections, results show that state-of-the-art image forgery methods do not transfer well to documents.

Fig. 5 visually presents the detection results for five representative forged documents from the SUPATLAN-TIQUE [43] dataset: `s1_73_c`, `s1_71_c`, `s11_41_c`, `s1_70_c` and `s11_34_c`. Each column corresponds to one document. The first row shows the forged image with manipulated characters highlighted in red. The second and third row display the outputs of the proposed noise-based and chromatic-based detection methods, respectively. The fourth row shows the best result obtained by an existing image forgery detection method, chosen individually for each document based on the highest $F1$ -score achieved.

For documents `s1_73_c` and `s1_71_c`, the results obtained by both methods are the same and match the manipulated zones of the document. This is because most inserted characters in scanned documents present both artifacts, as shown in Fig. 1. Conversely, document `s11_41_c` exhibits an example where the detection based on chromatic artifacts is able to spot the forgeries, while the noise-based method is not. This example confirms the analysis made in the previous paragraphs about the detection capabilities of both methods. Note that, in none of these examples, the best performing image forgery detection methods, which are FOCAL (documents `s1_73_c` and `s11_41_c`) and Adaptive (document `s1_71_c`), obtain a good prediction.

On document `s1_70_c` the results of both proposed methods are again similar, but we observe a few missed detections. Indeed, some manipulated words such as *né*, *à*, *au* and *Marc* are not detected. This fact reveals one of the flaws of the methods: short words are very hard (or even impossible) to detect. For example, in the case of a 2-character word, the variable Z in Eq. 6 follows a $\text{Bin}(2, 0.1)$ distribution. Even in the most favorable case of $K_\alpha(W) = 2$, we have $\mathcal{P}(Z \geq K_\alpha(W)) = 0.1^2$. With an NFA threshold of 0.01, it is impossible to detect such forgeries. On the other hand, we observe that the manipulated text *5100219* is partially detected. This is a consequence of the text mask given by the EAST detector. Indeed, in this case, the text mask separates *N°5100219* in two: *N°51* and *00219*. Finally, we observe the word *le* being detected together with the forged dates, even if it is not forged. This is again due to the text mask given by the EAST detector, which merges this word with the subsequent date. TruFor is the top-performing method, yielding some accurate predictions but also numerous false

alarms.

On document `s11_34_c` the chromatic method delivers false detections. Here, we observe that the falsely detected text presents a different color than the rest of the text present in the document. This characteristic is not taken into account by our method, as we assume that the chromatic artifacts are the same for all the characters in the text. However, chromatic artifacts depend on the components of the incoming light, which differ from one color to another [34]. It must also be highlighted that this color issue does not affect the noise method for which none of these false detections appear. In this case, Noisesniffer is the best-performing image forgery detection method even though its two detected regions do not match the forgeries.

5. Future work

The presented approaches are subject to several potential improvements and extensions. We only analyzed characters having lower noise levels than the rest of the characters. This step could be adapted to also include characters with higher noise level, a case that could be encountered in spliced documents. We assume a single model for all characters. This is not necessarily true for colored text, as shown in the previous section. Investigating the influence of the parameter α or testing for several values could be beneficial to avoid false negatives in short words. Finally, the proposed approaches could use a better word extractor.

A special mention needs to be made regarding binarized documents. This kind of document obviously falls out of the scope of the chromatic method. Still, the noise-based method could be adapted to them. Indeed, binarized documents still present salt and pepper noise, while newly typed characters should not present any noise at all. This anomaly could be detected with the current formulation of the method by using more sophisticated character extractors. Indeed, extracting the connected components leaves part of the noise in the background.

6. Conclusion

In this paper, we proposed two methods to detect digitally inserted characters in scanned documents. Our approaches rely on a statistical analysis of intensity fluctuations and chromatic artifacts found in characters. We show that both methods deliver high precision levels while keeping a good recall, especially the chromatic-based method. Such high-precision document forensics enables the automatic evaluation of massive data. We also include what is, to the best of our knowledge, the first evaluation of common image forgery detections in a document forgery dataset. The results of this benchmarking show a poor performance which should act as a call for the community to start developing methods that specifically focus on these types of images.

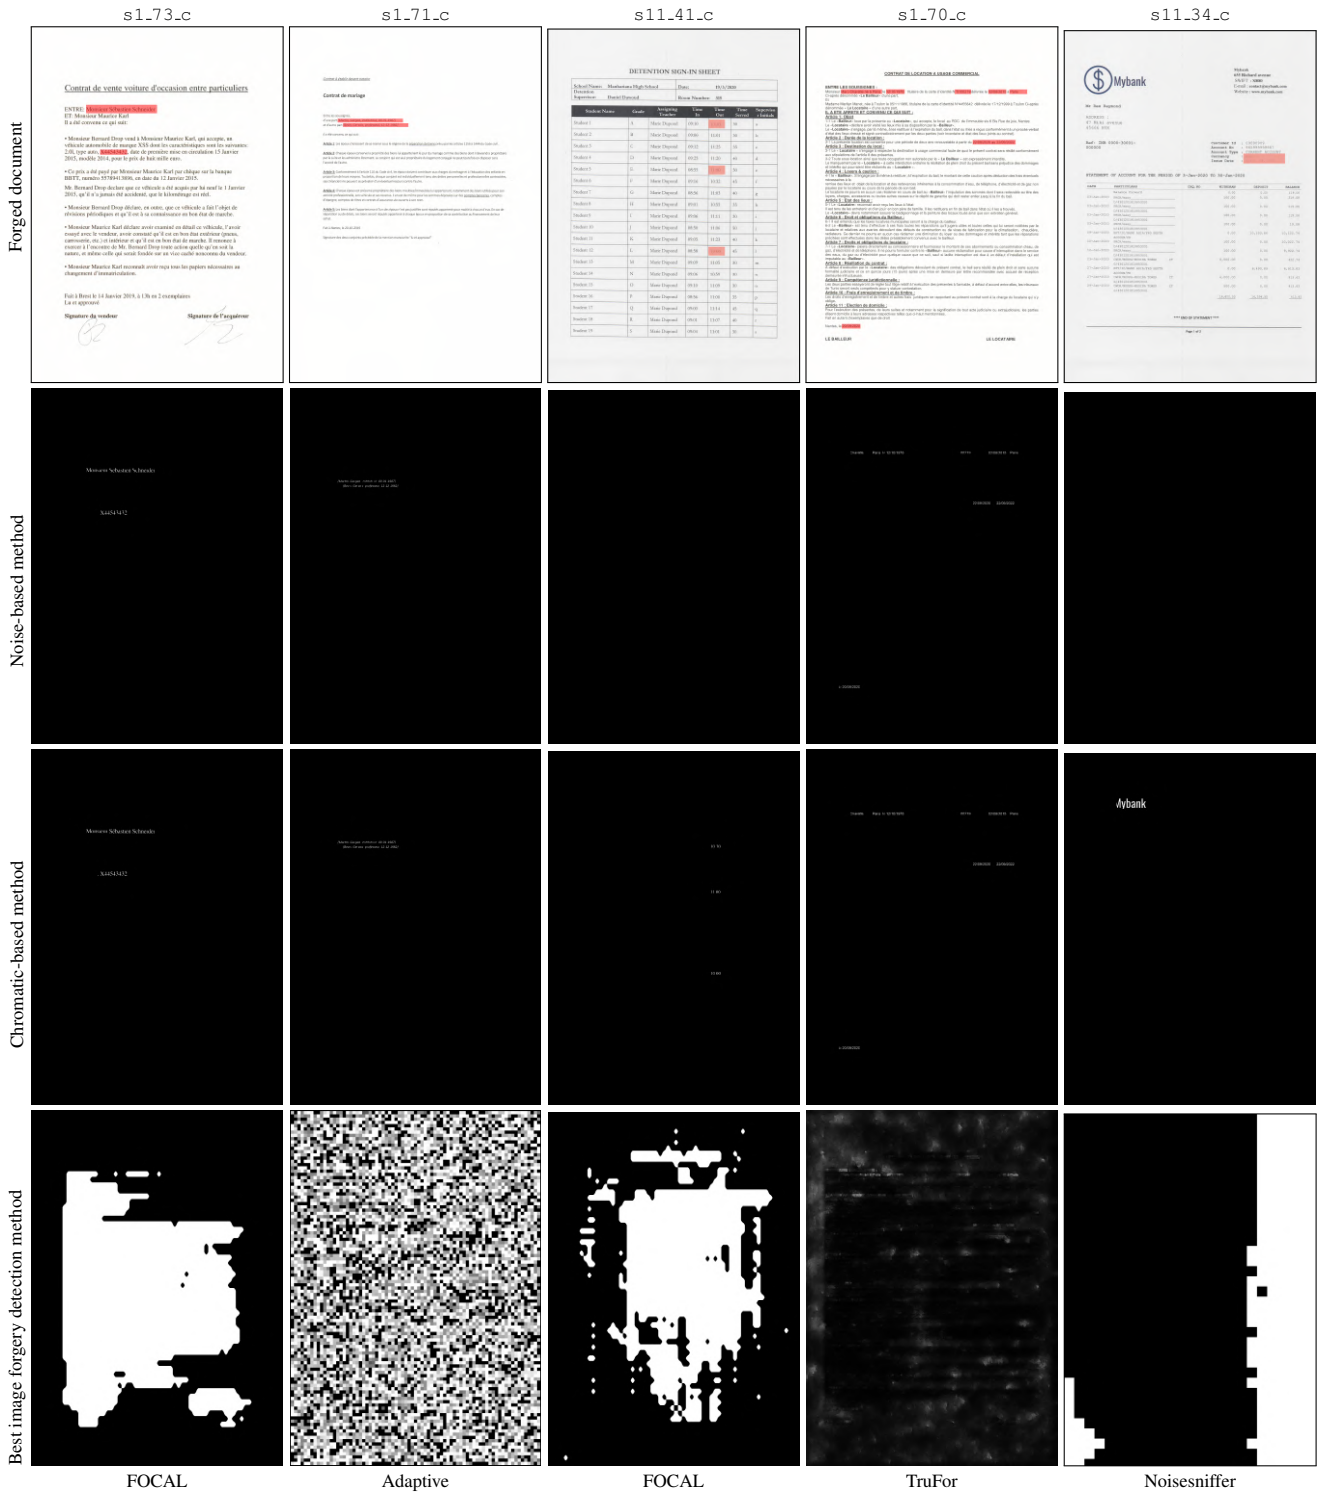


Figure 5. Results of the proposed approaches on documents s1_73_c, s1_71_c, s11_41_c, s1_70_c, and s11_34_c from the SUPAT-LANTIQUE [43] dataset. Each column corresponds to a forged document. Rows show, respectively, the forged document with forgeries annotated in red, the result of the noise-based method, the chromatic-based method, and the best-performing image forgery detection method for that example, in terms of $F1$ -score. Note that there are no good predictions made by the evaluated image forgery detection methods.

Acknowledgments

Work supported by Bpifrance, CollabNext project.

References

- [1] Chloé Artaud, Nicolas Sidère, Antoine Doucet, Jean-Marc Ogier, and Vincent Poulain D’Andecy Yooz. Find it! fraud detection contest report. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 13–18, 2018. 2
- [2] Fred Attneave. Some informational aspects of visual perception. *Psychological review*, 61(3):183, 1954. 4
- [3] Quentin Bamme, Rafael Grompone von Gioi, and Jean-Michel Morel. An adaptive neural network for unsupervised mosaic consistency analysis in image forensics. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14182–14192, 2020. 2, 3, 5, 6
- [4] Chaima Ben Rabah. *Analysis of scanned documents for integrity and authenticity checking*. Theses, Ecole nationale supérieure Mines-Télécom Atlantique ; École supérieure des communications de Tunis (Tunisie), 2021. 2
- [5] Romain Bertrand, Petra Gomez-Kramer, Oriol Ramos Terrades, Patrick Franco, and Jean-Marc Ogier. A System Based on Intrinsic Features for Fraudulent Document Detection. In *2013 12th International Conference on Document Analysis and Recognition*, pages 106–110, Washington, DC, USA, 2013. IEEE. 2
- [6] Romain Bertrand, Oriol Ramos Terrades, Petra Gomez-Kramer, Patrick Franco, and Jean-Marc Ogier. A Conditional Random Field model for font forgery detection. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*, pages 576–580, Tunis, Tunisia, 2015. IEEE. 3
- [7] André Braz, Maria López-López, and Carmen García-Ruiz. Raman spectroscopy for forensic analysis of inks in questioned documents. *Forensic Science International*, 232(1): 206–212, 2013. 2
- [8] Albert Berenguel Centeno, Oriol Ramos Terrades, Josep Lladós I Canet, and Cristina Canero Morales. Evaluation of Texture Descriptors for Validation of Counterfeit Documents. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, pages 1237–1242, Kyoto, 2017. IEEE. 3
- [9] Zhongxi Chen, Shen Chen, Taiping Yao, Ke Sun, Shouhong Ding, Xianming Lin, Liujuan Cao, and Rongrong Ji. Enhancing tampered text detection through frequency feature fusion and decomposition. In *European Conference on Computer Vision*, pages 200–217. Springer, 2024. 3
- [10] Pei-Ju Chiang, Nitin Khanna, Aravind K Mikkilineni, Maria V Ortiz Segovia, Sungjoo Suh, Jan P Allebach, George T-C Chiu, and Edward J Delp. Printer and scanner forensics. *IEEE Signal Processing Magazine*, 26(2):72–83, 2009. 1
- [11] Chang-Hee Choi, Jung-Ho Choi, and Heung-Kyu Lee. CFA pattern identification of digital cameras using intermediate value counting. In *Proceedings of the Thirteenth ACM Multimedia Workshop on Multimedia and Security*, page 21–26, New York, NY, USA, 2011. Association for Computing Machinery. 2
- [12] Mike Claybourn and M Ansell. Using Raman spectroscopy to solve crime: Inks, questioned documents and fraud. *Science & justice : journal of the Forensic Science Society*, 40:261–71, 2000. 2
- [13] Davide Cozzolino and Luisa Verdoliva. Noiseprint: A CNN-based camera model fingerprint. *IEEE Transactions on Information Forensics and Security*, 15:144–159, 2020. 2
- [14] Davide Cozzolino, Diego Gagnaniello, and Luisa Verdoliva. Image forgery detection through residual-based local descriptors and block-matching. *2014 IEEE International Conference on Image Processing (ICIP)*, pages 5297–5301, 2014. 2
- [15] Davide Cozzolino, Giovanni Poggi, and Luisa Verdoliva. Efficient dense-field copy-move forgery detection. *IEEE Transactions on Information Forensics and Security*, 10(11):2284–2297, 2015. 2
- [16] Davide Cozzolino, Giovanni Poggi, and Luisa Verdoliva. Splicebuster: A new blind image splicing detector. *Conference: IEEE Workshop on Information Forensics and Security*, 2015. 2, 3, 6
- [17] Francisco Cruz, Nicolas Sidere, Mickael Coustaty, Vincent Poulain D’Andecy, and Jean-Marc Ogier. Local Binary Patterns for Document Forgery Detection. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, pages 1223–1228, Kyoto, 2017. IEEE. 3
- [18] Agnès Desolneux, Lionel Moisan, and Jean-Michel Morel. *From Gestalt Theory to Image Analysis: A Probabilistic Approach*. Springer Publishing Company, Incorporated, 1st edition, 2007. 4, 5
- [19] Li Dong, Weipeng Liang, and Rangding Wang. Robust text image tampering localization via forgery traces enhancement and multiscale attention. *IEEE Transactions on Consumer Electronics*, 2024. 3
- [20] Hany Farid. *Photo forensics*. MIT press, 2016. 2
- [21] Marina Gardella, Pablo Musé, Miguel Colom, and Jean-Michel Morel. Image forgery detection based on noise inspection: Analysis and refinement of the noisefilter method. *Image Processing On Line*, 14:86–115, 2024. <https://doi.org/10.5201/ipo1.2024.462>. 2, 3, 5, 6
- [22] Douglas Goltz, Michael Attas, Gregory Young, Edward Cloutis, and Maria Bedynski. Assessing stains on historical documents using hyperspectral imaging. *Journal of Cultural Heritage*, 11(1):19–26, 2010. 2
- [23] Fabrizio Guillaro, Davide Cozzolino, Avneesh Sud, Nicholas Dufour, and Luisa Verdoliva. Trufor: Leveraging all-round clues for trustworthy image forgery detection and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20606–20615, 2023. 2, 3, 5, 6
- [24] Gail Hodge and Nikkia Anderson. Formats for digital preservation: A review of alternatives and issues. *Information Services and Use*, 27:45–63, 2007. 2
- [25] Hailey Joren, Otkrist Gupta, and Dan Raviv. OCR Graph Features for Manipulation Detection in Documents, 2020. arXiv:2009.05158 [cs]. 3
- [26] Myung-Joon Kwon, In-Jae Yu, Seung-Hun Nam, and Heung-Kyu Lee. CAT-Net: Compression artifact tracing network

- for detection and localization of image splicing. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 375–384, 2021. 2, 3, 5, 6
- [27] Chris Lennard, Moteaa M El-Defdar, and James Robertson. Forensic application of laser-induced breakdown spectroscopy for the discrimination of questioned documents. *Forensic science international*, 254:68–79, 2015. 2
- [28] Li Li, Yu Bai, Shanqing Zhang, and Mahmoud Emam. Document forgery detection based on spatial-frequency and multi-scale feature network. *Journal of Visual Communication and Image Representation*, 107:104393, 2025. 3
- [29] Zhouchen Lin, Junfeng He, Xiaoou Tang, and Chi-Keung Tang. Fast, automatic and fine-grained tampered JPEG image detection via DCT coefficient analysis. *Pattern Recognition*, 2009. 2, 3, 6
- [30] Xiaohong Liu, Yaojie Liu, Jun Chen, and Xiaoming Liu. Pscnet: Progressive spatio-channel correlation network for image manipulation detection and localization. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(11):7505–7517, 2022. 2, 3, 5, 6
- [31] Dongliang Luo, Yuliang Liu, Rui Yang, Xianjin Liu, Jishen Zeng, Yu Zhou, and Xiang Bai. Toward real text manipulation detection: New dataset and new solution. *Pattern Recognition*, 157:110828, 2025. 3
- [32] Siwei Lyu, Xunyu Pan, and Xing Zhang. Exposing region splicing forgeries with blind local noise estimation. *International Journal of Computer Vision*, 110:202–221, 2013. 2
- [33] Babak Mahdian and Stanislav Saic. Using noise inconsistencies for blind image forensics. *Image and Vision Computing*, 27:1497–1503, 2009. 2
- [34] David Marimont and Brian Wandell. Matching color images: the effects of axial chromatic aberration. *Journal of The Optical Society of America A-optics Image Science and Vision - J OPT SOC AM A-OPT IMAGE SCI*, 11, 1994. 7
- [35] Jose Melendez-Perez, Deleon Correa, Vinicius Hernandes, Damila Morais, Rodrigo Oliveira, Wanderley Souza, Jandyson Santos, and Marcos Eberlin. Forensic application of X-ray fluorescence spectroscopy for the discrimination of authentic and counterfeit revenue stamps. *Applied Spectroscopy*, 70, 2016. 2
- [36] Lokesh Nandanwar, Palaiahnakote Shivakumara, Umapada Pal, Tong Lu, Daniel Lopresti, Bhagesh Seraogi, and Bidyut B Chaudhuri. A new method for detecting altered text in document images. *International Journal of Pattern Recognition and Artificial Intelligence*, 35(12):2160010, 2021. 3
- [37] Anh D Nguyen, Hye-Young Kim, and Hoa N Nguyen. Taliu: A novel decoder and augmentation strategy for boosting tampered document image detection. *IEEE Access*, 2025. 3
- [38] Tina Nikoukhah, Jérémy Anger, Thibaud Ehret, Miguel Colom, Jean-Michel Morel, and Rafael Grompone von Gioi. JPEG grid detection based on the number of DCT zeros and its application to automatic and localized forgery detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 110–118, 2019. 2, 3, 5, 6
- [39] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979. 5
- [40] Julián O’Flaherty, Rodrigo Paganini, Juan Pablo Sotelo, Julieta Umpiérrez, Marina Gardella, Matías Tailanián, and Pablo Musé. Photoholmes: a python library for forgery detection in digital images. *Multimedia Tools and Applications*, pages 1–25, 2025. 6
- [41] Sofia Pessanha, Marta Manso, Ana Guilherme Buzanich, Mário Costa, and Maria Carvalho. Investigation on historical documents for forensic purposes by X-ray fluorescence spectrometry. *Surface and Interface Analysis*, 42:419, 2010. 2
- [42] Chenfan Qu, Chongyu Liu, Yuliang Liu, Xinhong Chen, Dezhi Peng, Fengjun Guo, and Lianwen Jin. Towards robust tampered text detection in document image: New dataset and new solution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5937–5946, 2023. 3
- [43] Chaima Ben Rabah, Gouenou Coatrieux, and Riadh Abdelfattah. The supatlantique scanned documents database for digital image forensics purposes. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 2096–2100, 2020. 3, 6, 7, 8
- [44] Abderrahmane Rahiche and Mohamed Cheriet. Forgery detection in hyperspectral document images using graph orthogonal nonnegative matrix factorization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2823–2831, 2020. 2
- [45] Hyun Jun Shin, Jong Ju Jeon, and Il Kyu Eom. Color filter array pattern identification using variance of color difference image. *Journal of Electronic Imaging*, 26(4):043015, 2017. 2
- [46] Carolina S. Silva, Maria Fernanda Pimentel, Ricardo S. Honorato, Celio Pasquini, José M. Prats-Montalbán, and Alberto Ferrer. Near infrared hyperspectral imaging for forensic analysis of document forgery. *Analyst*, 139:5176–5184, 2014. 2
- [47] Panagiotis C Theofanopoulos and Georgios C Trichopoulos. A novel fingerprint scanning method using terahertz imaging. In *2018 IEEE International Symposium on Antennas and Propagation & USNC/URSI National Radio Science Meeting*, pages 2463–2464. IEEE, 2018. 2
- [48] Beatriz Martínez Tornés, Emanuela Boros, Antoine Doucet, Petra Gomez-Krämer, Jean-Marc Ogier, and Vincent Poulain d’Andecy. Knowledge-based techniques for document fraud detection: a comprehensive study. In *International Conference on Computational Linguistics and Intelligent Text Processing*, pages 17–33. Springer, 2019. 3
- [49] Beatriz Martínez Tornés, Emanuela Boros, Antoine Doucet, Petra Gomez-Krämer, and Jean-Marc Ogier. Detecting forged receipts with domain-specific ontology-based entities & relations. In *International Conference on Document Analysis and Recognition*, pages 184–199. Springer, 2023. 3
- [50] Haiwei Wu, Yiming Chen, and Jiantao Zhou. Rethinking image forgery detection via contrastive learning and unsupervised clustering, 2023. 2, 3, 5, 6
- [51] Chenhao Zheng, Ayush Shrivastava, and Andrew Owens. Exif as language: Learning cross-modal associations between im-

ages and camera metadata. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6945–6956, 2023. [2](#), [3](#), [5](#), [6](#)

- [52] Xinyu Zhou, Cong Yao, He Wen, Yuzhi Wang, Shuchang Zhou, Weiran He, and Jiajun Liang. East: An efficient and accurate scene text detector. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2642–2651, 2017. [5](#)