JADE OWL: JPEG 2000 forensics by wavelet offset consistency analysis

Quentin Bammey

Université Paris-Saclay, ENS Paris-Saclay, CNRS, Centre Borelli

Gif-sur-Yvette, France

ORCID: 0000-0003-2280-2349

Abstract—In a world teeming with digital images, the credibility of visual data has become of paramount importance. While it is now simpler than ever to manipulate an image for malicious purposes such as misinformation, the tools for detecting such alterations have predominantly been developed for either uncompressed or JPEG-compressed natural images. However, medical and satellite imagery, domains where the potential for fraud is high, often use a different compression format – JPEG 2000. We present a JPEG 2000 Anomaly Detection Estimator via Offset of Wavelet Localization – the Jade Owl –, a novel method for detecting forgeries in JPEG 2000 images by analyzing the consistency of traces left by its compression. Our technique hinges on the premise that the wavelet coefficients of a JPEG 2000 image are lower when the same offset is applied during the wavelet transform than they are when the offset is different. By employing this principle locally, we're able to detect regions with significantly different offsets, indicating potential forgeries such as copy-move. An accompanying *a contrario* model further refines this detection to make automatic detections while controlling false positives. To evaluate the method, we've created a unique dataset of JPEG 2000 forgeries. This novel approach significantly paves the way for JPEG 2000 image forensics, introducing a sensitive and efficient tool for authenticity verification in critical sectors such as healthcare and satellite imagery.

Index Terms—Image forgery detection, JPEG 2000, a contrario detection, medical image forensics, wavelet analysis

I. INTRODUCTION

The trustworthiness of images has become a pivotal concern in our increasingly digital world. With image manipulation more accessible than ever, the ability to verify the authenticity of visual data is of paramount importance. While the majority of image forgeries stem from altered photographs, the methods to detect these have predominantly focused on traces left by the image signal processing pipeline and JPEG compression artifacts. However, images compressed with JPEG 2000, a compression standard often employed in the realms of medical and satellite imagery, present a unique challenge. While the use of this compression standard is limited outside of these two domains, both feature a high risk of forgeries. Medical images can be forged for falsely advising treatment, or to protect oneself from lawsuits by pretending something of which a physician should have been aware was actually impossible to detect on the scans [1], [2]. Satellite images

(a) Original JPEG 2000-compressed image

(b) Multilevel wavelet representation of the image

(c) Same, but the image is shifted by one pixel in both direction before the wavelet tranform

Fig. 1: Toy example: On the left, an image is JPEG 2000-compressed. On the middle, the wavelet representation of the compressed image is shown in absolute values. On the right, the image is shifted by one pixel in both direction before computing the wavelet representation. This simple shift causes the representation to contain much higher values in magnitude. This can be used to retrieve inconsistencies in JPEG-2000 images: computing the wavelet transform help finds the original offset of the image, which can is locally disturbed by many forgeries such as copy-move, enabling one to detect such forgeries by locally estimating the offset.

can be manipulated even more maliciously by governments or large non-governmental organizations to lie about the presence or absence of nuclear or military equipment and facilities, either to avoid drawing attention, to wrongfully incriminate other countries, or to spread disinformation during war [3], [4].

JPEG 2000 compression operates in the wavelet domain, performing a multi-scale wavelet transform followed by quantization of wavelet coefficients. This process inherently leaves distinctive traces in the compressed image. Those fragile traces are easily disturbed by post-compression alterations. As image forgeries imply localized alterations, local inconsistencies in the detected traces are clear proofs of image forgeries. Even a simple copy-move forgery within the same image can offset values and disrupt quantization. Similarly, the localized absence of quantization in a region of a JPEG 2000 compressed image could indicate forgery by splicing, where a donor object would be taken from an uncompressed or standard JPEG-compressed image and pasted onto the target JPEG-2000

image.

In this paper, we introduce a method that harnesses these characteristics to detect inconsistencies arising from JPEG 2000 compression and locate forgeries. This method exploits the property that the wavelet transform coefficients of a JPEG 2000 image are lower when the image is recompressed with the same offset, as visualized in Figure 1. By conducting this process locally, our method can identify regions with significant offset variance, indicating potential forgery. Additionally, an a contrario detection model is implemented to further refine this detection, ensuring a controlled false positive rate.

We validate our approach using a novel dataset of forged JPEG 2000 images, specifically created for this purpose. Our method demonstrates a high sensitivity to JPEG 2000 compression inconsistencies, outperforming existing forensic tools in this regard. This approach expands the frontier of JPEG 2000 image forensics, providing a robust tool for forgery detection in sectors such as medical and satellite imaging where the risk of fraudulent modification is significant.

II. RELATED WORKS

Image forensics in photographic images has been the focus of extensive research these past few years. Forgery detection methods are usually divided into three categories. The first, older category consists of methods consists in analysing the high-frequency traces in an image to find inconsistencies in it, including but not limited to demosaicing traces [5], [6], noise level analysis [7] or JPEG compression artefacts [8].

A newer category of methods consist of neural networks trained directly on forged [9] or authentic images to analyse their noise patterns and detect anomalies [10]. This approach has already been extended onto satellite images. The authors of SatSVDD [4] propose a deep learning-based method for detecting and localizing splicing manipulations in satellite images. SatSVDD utilizes an autoencoder and a support vector data description classifier to capture the salient features of pristine images and distinguish them from forged patches. The authors of SeaTheSeams use a two-stage approach to detect and locate seam carving manipulations in satellite imagery. In the first stage, a fully convolutional network is employed to localize seams, and in the second stage, a convolutional neural network performs binary classification to detect seam carving manipulations. The authors of SatSVDD [4] propose a deep learning-based method for detecting and localizing splicing manipulations in satellite images. SatSVDD utilizes an autoencoder and a support vector data description classifier to capture the salient features of pristine images and distinguish them from forged patches. The authors of SeaTheSeams use a two-stage approach to detect and locate seam carving manipulations in satellite imagery. In the first stage, a fully convolutional network is employed to localize seams, and in the second stage, a convolutional neural network performs binary classification to detect seam carving manipulations.

Some research has also already been focused on medical images. Ghoneim et al. [1] propose a medical image forgery detection system for the healthcare framework, utilizing a noise map and regression filter for analysis. The system ensures the integrity of healthcare-related images by employing support-vector-machine-based and extreme-learning-based classifiers. Manimurugan and Porkumaran [11] introduce a secure visual cryptography algorithm for black and white medical images. It divides the images into two shares, reconstructs the original image without post-processing, and incorporates forgery detection using JPEG Ghosts [12] for enhanced security. The proposed approach enables secure storage and compressed transmission of medical images while ensuring data integrity. However, it focuses on JPEG artefacts, whereas medical images are mostly JPEG-2000-compressed. Ulutas et al [13] focus on the integrity protection of medical images transmitted over the Internet for telemedicine purposes. The proposed method utilizes passive image authentication using local binary pattern rotation invariant (LBPROT) and scale invariant feature transform (SIFT) [14] to detect tampered regions in medical images without the need for watermarking or extra information embedding, while keeping robustness to attacks and postprocessing such as scaling or rotation.

Although some research has already been focused specifically on medical or satellite images, none target the JPEG 2000 artefacts. Recent works mainly propose directly analysing images with generic neural networks, or detect specific traces such as seam carving.

Existing works on JPEG 2000 image forensics are scarce. Qadir et al. [15] propose the use of Benford's Law to estimate JPEG2000 compression rates in image forensics. By analyzing the 1st digit probabilities of DWT coefficients, the compression rates can be determined. The method aims to provide a reliable approach for detecting and estimating JPEG2000 compression in images. Zhang et al. [16] present a new approach for detecting tampered images that have undergone double JPEG 2000 compression. The method analyzes the statistics and artifacts in the Fourier transforms of DWT coefficient histogram to identify the singularity in mid and high frequencies, indicating double compression. Thresholding is then applied to detect double JPEG 2000 compression, and the percent of non-zero DWT coefficients in high frequencies is examined to locate tampered areas.

The *a contrario* detection theory, as detailed in the works of Desolneux et al. [17], [18], offers a valuable framework to overcome this challenge. The crux of this theory is rooted in the concept of non-accidentalness. It proposes the detection of data that deviate from the expected norm under a background hypothesis, thereby rendering them statistically improbable or non-accidental. The *a contrario* method effectively thresholds results by setting a tolerated limit on the number of false alarms (also known as NFA - Number of False Alarms) that one is willing to accept under the prescribed hypothesis. This approach to detection has found successful application across a broad spectrum of detection tasks, as seen in numerous studies [19]–[37] that showcase the theory's effective utilization in various facets of image processing. This includes but is not limited to image forensics, where the *a contrario* method has proven to be of significant value [6]–[8], [38]–[44]. In

the case of image forensics, this approach bridges a critical gap left by many existing forgery detection methods. The *a contrario* theory has thus emerged as a significant cornerstone in modern forgery detection processes, shedding light on potentially doctored areas and enabling more reliable and efficient detection of image manipulation.

III. PROPOSED METHOD

The proposed method is an image forensics tool designed to detect forgeries in JPEG 2000 images. It achieves this by leveraging the shift variance property inherent to the wavelet transformation applied during JPEG 2000 compression process.

The algorithm exploits a particular characteristic of this compression process: if the wavelet transform is applied once more to the compressed image using the same offset as in the original compression, the retrieved wavelet coefficients will be similar to the quantized coefficients from the original compression. However, if there is a change in offset, the wavelet coefficients will differ markedly from the original set, showing greater variability, and more importantly, an increased number of large amplitude coefficients.

Given a one-channel image, we compute their wavelet transform with the Cohen–Daubechies–Feauveau 9/7 wavelet, which is usually used for lossy JPEG 2000 compression[1].

Although the JPEG 2000 wavelet transform usually contains five scales, we only analyse the finest two. Indeed, our experiments showed finding forgeries at coarser levels was usually impossible. The two levels are checked separately, then the resulting forgery masks for both are merged. From now on, we thus assume we are only analysing the finest or the second-finest scale.

The first major step of the method is the offset estimation. The number of possible offsets in each direction is 2^{s+1}, where s is the scale (0 is the finest scale). The finest level can have a 0 or 1 offset, the second finest an offset between 0 and 3. For each possible pair of offsets (horizontal and vertical offset), we compute the multilevel wavelet transform with said offset and look at the desired level. At each pixel (one pixel having a different resolution depending on the scale we consider), we estimate locally the best offset pair as the one whose coefficients are the smallest in magnitude. The horizontal offset is computed separately using the horizontal and diagonal coefficients, the vertical offset using the vertical and diagonal coefficients. If the two maps disagree, the estimation is considered locally invalid at this location.

Of course, such an estimation is unreliable at the pixel level. Aggregating the results is necessary to provide reliable information.

Each pixel thus "votes" for one location. The globally best phase is the one that receives the most votes across the image.

The next step in the method is the local offset estimation. This involves checking smaller windows of the image instead of the entire image at once. Given a window size of W, the algorithm moves a window over the image and for each window, performs the same process as the global estimation to determine the offset that minimizes the number of large amplitude coefficients. The result is a map of local offsets for the image, indicating the most likely compression offset for each local region.

The third and final step in the process is the forgery detection. For each local window, the algorithm compares the local offset to the global offset. Intuitively, one could think that regions where the local offset is different than the global offset are forged. However, the estimation might simply be noisy and the difference between the local and global offset might be due to chance. Indeed, the pixelwise comparison used to check each pixel offset is hardly perfect, and is only accurate by aggregating results. In an extreme case, where the image had no traces of JPEG 2000 compression whatsoever, the pixels estimations would be uniformly random, and one of the four offsets would still have more votes than the other ones, both globally and locally. Three quarters of the windows would thus be considered forged.

This is where *a contrario* detection comes to play. In the global image, let us suppose that n pixels have provided a valid vote, out of which k were for the most likely offset, $\delta^\star$ (the other three offsets thus have less than k votes each). Under a background hypothesis where the image was not JPEG-2000 compressed, the votes are expected to be uniformly random. The likelihood that $\delta^\star$ received k or more votes is thus the tail of a binomial,

$$\mathbb{P}\left(n_{\text{votes}} \geq k\right) = \sum_{i=j}^{n} \binom{n}{j} \cdot \left(\frac{1}{4}\right)^j \left(\frac{3}{4}\right)^{n-j}. \quad (1)$$

As there are four possible offsets, the probability that any of the offset received at least k votes is four time this value. It represents the number of false alarms (NFA) associated with the global detection:

$$\mathbb{N}FA_{\text{global}} = 4 \cdot \sum_{i=j}^{n} \binom{n}{j} \cdot \left(\frac{1}{4}\right)^j \left(\frac{3}{4}\right)^{n-j}. \quad (2)$$

Thresholding on this value can thus be done simply by selecting the frequency at which one would accept to see a (false) detection on uncompressed images. For instance, setting the threshold at 10^{-3} ensures that, under the background hypothesis of an image having no traces of JPEG 2000 compression, at most one every thousand images will be wrongly detected as having a best offset. If the NFA of the detection is over the set threshold[2], we cannot conclude at the presence of JPEG 2000 compression traces, and the method stops there (at least for this scale) without finding forgeries.

The principle is similar to detect local inconsistencies. However, what we want to limit is the frequency at which

---

[1]Of course, the method can be used with other wavelets, should there be a doubt about which specific transform was used. In a similar way to finding the best offset, it is possible to detect the wavelet used for the compression, as it will show signs of quantization.

[2]Note that significant event have a low NFA, close to zero, while insignificant events have a high NFA.

images are wrongly detected as forged, not the frequency at which a window is wrongly detected. To provide an upper bound on the tolerated number of false alarms, the detection NFA is thus multiplied by the number of independent windows in an image, as well as by the number of scales and the number of window sizes used. Each window is thus considered forged if the two conditions are met:

- The most-voted offset in this window is not the globally-detected one, and
- The NFA associated with the detection is below the set threshold.

Assuming a NFA threshold of 10^{-3}, this ensures that under the background hypothesis, at most one every thousand images will have inconsistent windows detected. Windows with an NFA below a predefined threshold are flagged as detected forgeries.

This whole process is conducted iteratively for each window size, and for at the two finest scales of the wavelet decomposition to capture any potential forgeries at various scales and orientations. To keep consistency over the tolerated number of false alarms, both the global and local NFA scores are thus multiplied by two (the number of scales) before checking against the threshold. The local NFA scores are further multiplied by the number of studied window sizes.

The output is a binary forgery detection map that shows the locations of potential forgeries within the image.

Until now, we have assumed we were processing a one-channel, black-and-white image. If a colour image is analysed instead, it is first converted to the YCbCr colour space, the colour space in which the wavelet transform is usually performed during JPEG 2000 compression. Each channel is then analysed separately for forgeries. In this case, the NFA are again multiplied by three before checking them against a threshold.

By exploiting the shift variance property of the wavelet transform, this method offers a new approach to detecting forgeries in JPEG 2000 images, potentially strengthening the security and reliability of images in areas such as medical imaging and satellite imagery.

IV. DATASET AND EXPERIMENTS

To robustly assess the effectiveness of our novel methodology, we require a dataset composed of forged JPEG2000 images. We generate a new dataset, deriving inspiration from the methodology employed in the creation of the Trace database [48].

The procedure to create forgeries involves compressing each chosen image twice, each time with a different offset. These two versions of the image are subsequently amalgamated following a predetermined forgery mask. The resulting composite images are visually indistinguishable from the original, unaltered counterparts. Yet, upon closer examination, these forged images betray inconsistencies in their compression patterns; a feature that enables us to test both our proposed method and other state-of-the-art algorithms for their proficiency in identifying forgeries via JPEG 2000 compression traces.

Method	MCC	F1
Proposed (both levels)	**0.540**	**0.582**
Proposed (level 0 only)	0.489	0.517
Proposed (level 1 only)	0.509	0.534
ManTraNet [9], [45]	-0.006	0.189
Noisesniffer [7]	-0.009	0.035
Noiseprint [10]	0.006	0.080
Comprint [46]	0.005	0.070
Splicebuster [47]	-0.005	0.205
AdaCFA [5]	0.178	0.324
ZERO [8]	0.000	0.000

TABLE I: Evaluation of the proposed method and publicly available state-of-the-art on the Trace-J2k dataset. The proposed method is to date the only one able to spot inconsistencies in the JPEG 2000 compression traces.

Levels used	16 bits	8 bits	JPEG $Q = 98$	JPEG $Q = 95$
Both levels	0.540	0.253	0.125	0.075
Level 0 only	0.489	0.146	0.064	0.036
Level 1 only	0.509	0.236	0.111	0.067

TABLE II: MCC of the method whan the image loses precision to 8 bits and is JPEG-compressed, as could happen in the wild. As can be seen here, the proposed method can still work on 8-bits images, and even against a small JPEG compression, although its performances are expectedly diminished.

Our source images and forgery masks align with those used in the original Trace database, which themselves are derived from the Raise raw images dataset [49]. We choose to keep using photographic images rather than specific medial or satellite images so as to provide generalizable results, as the JPEG 2000 are not specific to the kind of image.

For the purpose of comparison, we measure our method's performance against several well-known, publicly available SOTA techniques. These include Splice-buster [47], Noiseprint [10], Noisesniffer [7], Comprint [46], ManTraNet [9], [45], AdaCFA [5] and ZERO [8], [43]. We quantify the results using Matthew's Correlation Coefficient (MCC), a metric that varies from -1 to 1; here, a score of 1 signifies a flawless detection, -1 its inverse, and 0 suggests an absence of correlation between the detection results and the ground truth. Additionally, we measure the F1 score, which provides a balanced measure of precision and recall.

As seen in Table I and visually illustrated in Figure 2, our proposed method is the only one capable of reliably spotting inconsistencies within the JPEG 2000 compression. The only other method showing some level of sensitivity to these discrepancies is AdaCFA [5], an unexpected revelation given that this technique focuses primarily on identifying traces of demosaicing. The proposed method remains robust when the image loses in precision and becomes an 8-bits image, and even against small JPEG compression after the forgery is done, although its performance is expectedly worse.

In Table III, we analyse the results of the method using the finest scale only, the second-finest scale only, and both scales, depending on the actual offset of the forgery modulo 2. We
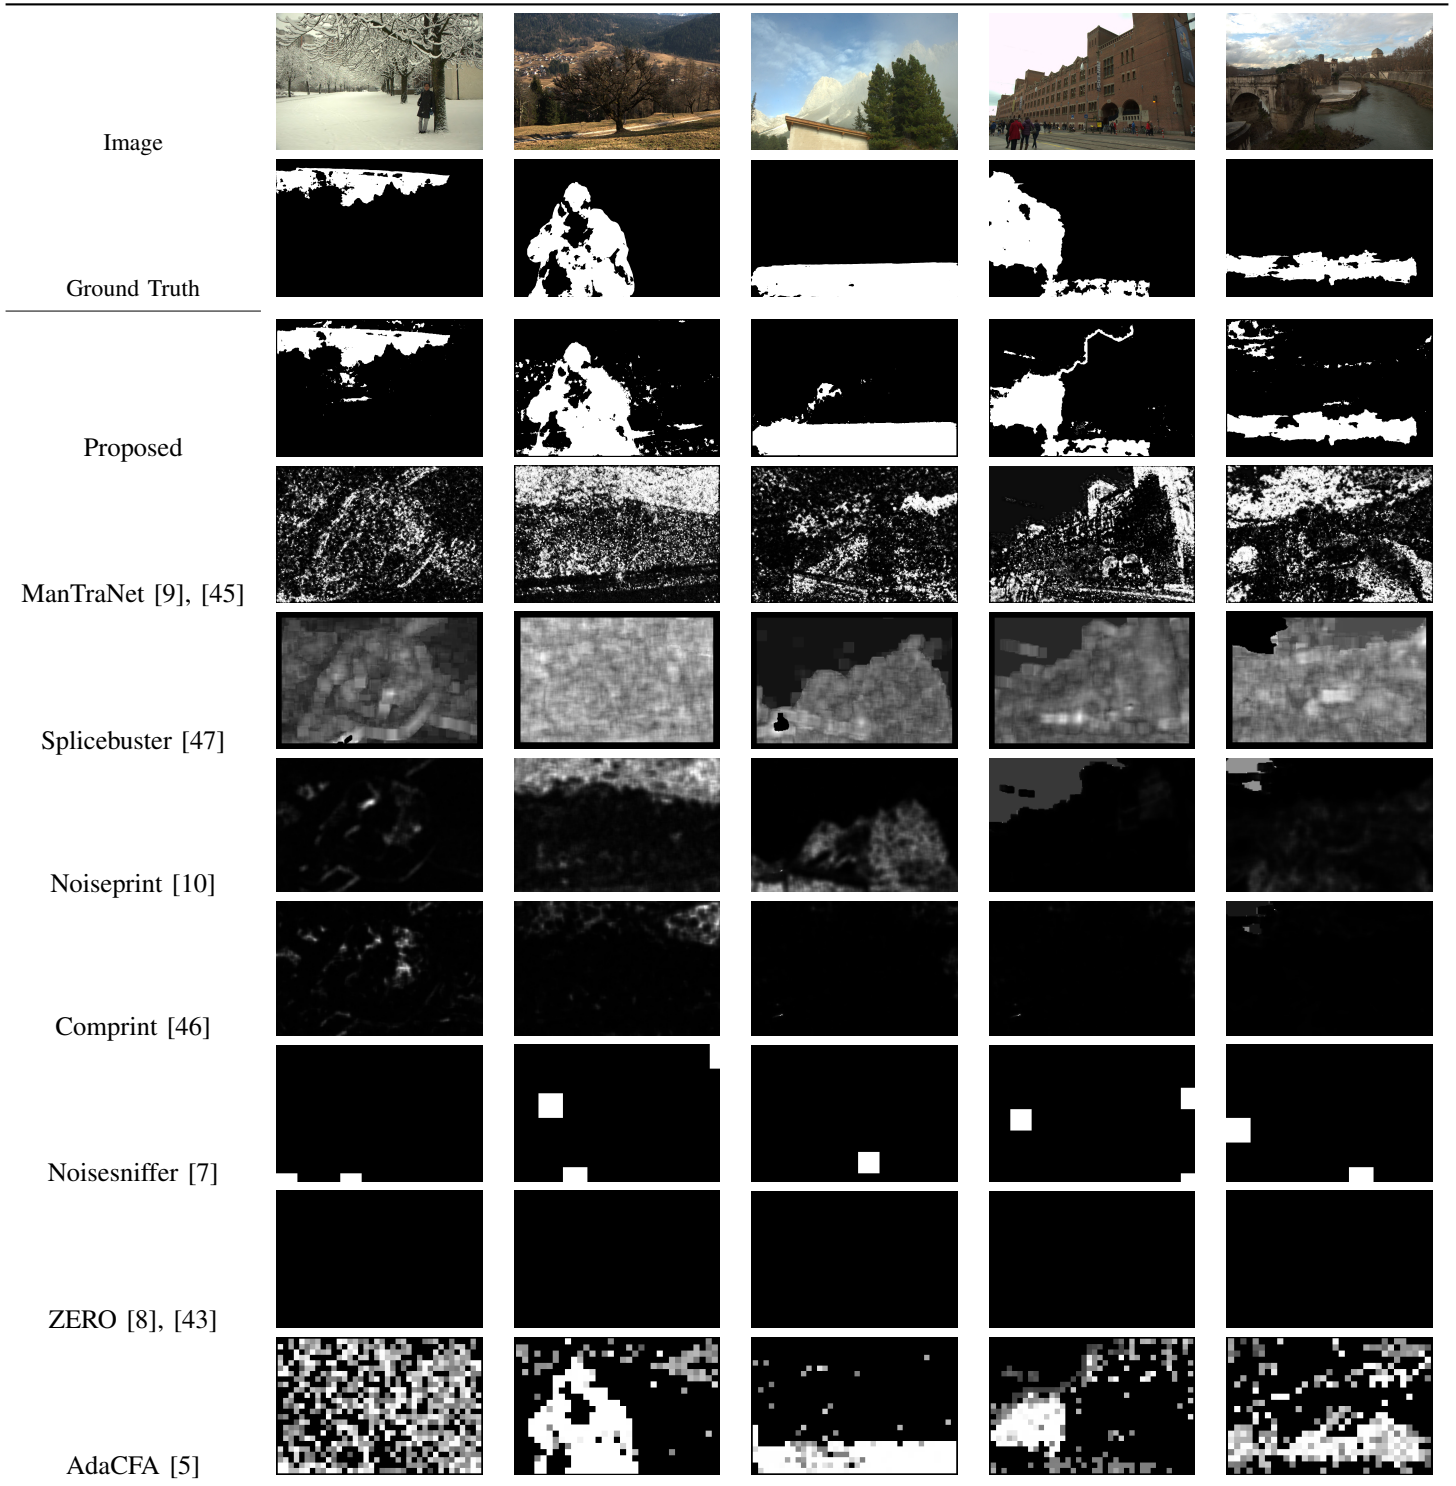
Fig. 2: Visual results on the proposed dataset. The proposed method globally localizes forgeries well enough, although there are several false positives throughout the images. Surprisingly, AdaCFA is also able to detect inconsistencies despite focusing on another kind of traces. Other methods, including generic ones such as MantraNet, Splicebuster and Noiseprint, are not sensitive at all to JPEG 2000 compression inconsistencies.

	$\delta_x \equiv 0$	$\delta_x \equiv 1$
$\delta_y \equiv 0$	0.00 – **0.44** – 0.44	**0.63** – 0.54 – 0.59
$\delta_y \equiv 1$	**0.69** – 0.54 – 0.59	**0.66** – 0.52 – 0.54

TABLE III: Results of the method using the finest scale only – the second-finest scale only – and both scales, depending on the actual offset of the forgery modulo 2. We can see that looking at the finest scale usually gives the best results when the falsification is not at a zero-offset modulo 2 in both directions. Indeed, in this case, no inconsistencies are present at the finest scale. On the other hand, the second-finest scale has a larger periodicity of 4, and will thus be able to make detections as long as the offset is not zero modulo 4 in both directions. As a consequence, looking at both scales provides the best results, as the second-finest scales improves the detection on the zero-modulo-2-offset cases.

can see that looking at the finest scale usually gives the best results when the falsification is not at a zero-offset modulo 2 in both directions. Indeed, in this case, no inconsistencies are present at the finest scale. On the other hand, the second-finest scale has a larger periodicity of 4, and will thus be able to make detections as long as the offset is not zero modulo 4 in both directions. As a consequence, looking at both scales provides the best results, as the second-finest scales improves the detection on the zero-modulo-2-offset cases.

V. DISCUSSION AND CONCLUSION

The compelling results of this study, despite the simplicity of the Jade owl method, reveal a significant gap in current forgery detection methods, particularly in the context of JPEG 2000 images, which are routinely used in domains of high-risk potential such as medical imaging and satellite imagery. Our novel method, Jade Owl, bridges this gap by leveraging the properties of wavelet transform and shift inconsistencies, which provides a new perspective for forgery detection. The capability of our approach to accurately detect forgeries, even in the absence of visible differences, has broad implications in safeguarding the integrity of visual data.

An unexpected finding in our comparative analysis was the performance of the AdaCFA method, which showed some sensitivity to JPEG 2000 compression inconsistencies despite its main focus on demosaicing traces. This points to the potential overlap in the detection methods and hints at the existence of common underlying features in various types of forgeries.

The creation of a new dataset with forged JPEG2000 images has the potential to catalyse further research and advancements in this area. It provides a valuable resource not only for testing our method but also for benchmarking future techniques and innovations in the field of forgery detection.

In conclusion, our research has culminated in the development of a robust and reliable method for detecting forgeries in JPEG2000 images, which significantly outperforms state-of-the-art methods in the domain. The Jade Owl method provides an innovative solution to a pressing challenge, ensuring the integrity and authenticity of images in critical areas such as healthcare and satellite imagery.

Going forward, we believe that this work will pave the way for more comprehensive and sophisticated forgery detection methods in the presence of JPEG 2000 compression. We have shown that the compression itself can present detectable inconsistencies when the image is forged. In the specific case of satellite images, this method is a first step towards the understanding of artefacts and traces present in the signal. The JPEG 2000 compression occurs before the images are sent to the ground, where they are further processed in a manner which depends on the satellite constellation. To accurately detect forgeries using JPEG 2000 traces, it will thus be necessary to analyse the artefacts over such processing.

More generally, future works should also focus on better understanding and modelling the behaviour of existing image features and traces, such as noise, when the image undergoes JPEG 2000 compression, similarly to what has already been done for JPEG compression [50]–[56]. Going forward, we believe that this work will pave the way for more comprehensive and sophisticated forgery detection methods. Furthermore, the dataset we have developed will provide a sound platform for researchers to validate their novel approaches and methods in the ever-evolving fight against image forgery.

The Jade Owl method and the new dataset will be made publicly available to the research community, fostering further exploration and innovation in the realm of image forensics. We look forward to seeing how these resources will be utilized and enhanced by the collective effort of the forensic research community.

REFERENCES

[1] A. Ghoneim, G. Muhammad, S. U. Amin, and B. Gupta, "Medical image forgery detection for smart healthcare," *IEEE Communications Magazine*, vol. 56, no. 4, pp. 33–37, 2018.

[2] C.-R. Piao, D.-M. Woo, D.-C. Park, and S.-S. Han, "Medical image authentication using hash function and integer wavelet transform," in *2008 Congress on Image and Signal Processing*, vol. 1. IEEE, 2008, pp. 7–10.

[3] C. Gudavalli, E. Rosten, L. Nataraj, S. Chandrasekaran, and B. S. Manjunath, "Seetheseams: Localized detection of seam carving based image forgery in satellite imagery," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2022, pp. 1–11.

[4] J. Horváth, D. Güera, S. K. Yarlagadda, P. Bestagini, F. M. Zhu, S. Tubaro, and E. J. Delp, "Anomaly-based manipulation detection in satellite images," *networks*, vol. 29, p. 21, 2019.

[5] Q. Bammey, R. G. von Gioi, and J.-M. Morel, "An adaptive neural network for unsupervised mosaic consistency analysis in image forensics," in *CVPR*, 2020.

[6] ——, "Forgery detection by internal positional learning of demosaicing traces," in *WACV*, January 2022, pp. 328–338.

[7] M. Gardella, P. Musé, J.-M. Morel, and M. Colom, "Noisesniffer: a fully automatic image forgery detector based on noise analysis," in *IWBF*. IEEE, 2021.

[8] T. Nikoukhah, J. Anger, T. Ehret, M. Colom, J.-M. Morel, and R. Grompone von Gioi, "JPEG grid detection based on the number of DCT zeros and its application to automatic and localized forgery detection," in *CVPRW*, 2019.

[9] Y. Wu, W. AbdAlmageed, and P. Natarajan, "Mantra-net: Manipulation tracing network for detection and localization of image forgeries with anomalous features," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

[10] D. Cozzolino and L. Verdoliva, "Noiseprint: A cnn-based camera model fingerprint," *IEEE TIFS*, 2020.

[11] S. Manimurugan and K. Porkumaran, "A new fast and efficient visual cryptography scheme for medical images with forgery detection," in *2011 International Conference on Emerging Trends in Electrical and Computer Technology*, 2011, pp. 594–599.

[12] H. Farid, "Exposing digital forgeries from jpeg ghosts," *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 1, pp. 154–160, 2009.

[13] G. Ulutas, A. Ustubioglu, B. Ustubioglu, V. V Nabiyev, and M. Ulutas, "Medical image tamper detection based on passive image authentication," *Journal of digital imaging*, vol. 30, pp. 695–709, 2017.

[14] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *International journal of computer vision*, vol. 60, pp. 91–110, 2004.

[15] G. Qadir, X. Zhao, and A. T. Ho, "Estimating jpeg2000 compression for image forensics using benford's law," in *Optics, Photonics, and Digital Technologies for Multimedia Applications*, vol. 7723. SPIE, 2010, pp. 133–142.

[16] J. Zhang, H. Wang, and Y. Su, "Detection of double-compression in jpeg2000 images for application in image forensics." *Journal of multimedia*, vol. 4, no. 6, 2009.

[17] A. Desolneux, L. Moisan, and J. Morel, "From gestalt theory to image analysis. interdisciplinary applied mathematics, vol. 35," pp. 7–23, Oct 2007. [Online]. Available: https://doi.org/10.1023/A:1026593302236

[18] A. Desolneux, L. Moisan, and J.-M. Morel, "Meaningful alignments," *IJCV*, 2000.

[19] P. Moulon, P. Monasse, and R. Marlet, "Adaptive structure from motion with a contrario model estimation," in *Computer Vision – ACCV 2012*, K. M. Lee, Y. Matsushita, J. M. Rehg, and Z. Hu, Eds., Berlin, Heidelberg, 2013.

[20] A. Davy, T. Ehret, J.-M. Morel, and M. Delbracio, "Reducing anomaly detection in images to detection in noise," in *IEEE ICIP*, 2018.

[21] A. Robin, L. Moisan, and S. Le Hegarat-Mascle, "An a-contrario approach for subpixel change detection in satellite imagery," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, 2010.

[22] T. Ehret, A. Davy, M. Delbracio, and J.-M. Morel, "How to Reduce Anomaly Detection in Images to Anomaly Detection in Noise," *Image Processing On Line*, vol. 9, 2019.

[23] T. Dagobert, R. Grompone von Gioi, C. de Franchis, J.-M. Morel, and C. Hessel, "Cloud Detection by Luminance and Inter-band Parallax Analysis for Pushbroom Satellite Imagers," *Image Processing On Line*, vol. 10, 2020.

[24] L. Moisan, P. Moulon, and P. Monasse, "Automatic Homographic Registration of a Pair of Images, with A Contrario Elimination of Outliers," *Image Processing On Line*, vol. 2, 2012.

[25] N. Mandroux, T. Dagobert, S. Drouyer, and R. Grompone von Gioi, "Single Date Wind Turbine Detection on Sentinel-2 Optical Images," *Image Processing On Line*, vol. 12, 2022.

[26] A. Buades, R. Grompone von Gioi, and J. Navarro, "Contours, Corners and T-Junctions Detection Algorithm," *Image Processing On Line*, vol. 8, 2018.

[27] C. Aguerrebere, P. Sprechmann, P. Muse, and R. Ferrando, "A-contrario localization of epileptogenic zones in spect images," in *2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, 2009.

[28] E. Aldea and S. L. Hégarat-Mascle, "Robust crack detection for unmanned aerial vehicles inspection in an a-contrario decision framework," *JEI*, 2015.

[29] R. Grompone von Gioi, C. Hessel, T. Dagobert, J.-M. Morel, and C. de Franchis, "Ground Visibility in Satellite Optical Time Series Based on A Contrario Local Image Matching," *Image Processing On Line*, vol. 11, 2021.

[30] R. Grompone von Gioi, J. Jakubowicz, J.-M. Morel, and G. Randall, "Lsd: A fast line segment detector with a false detection control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 32, 2010.

[31] ——, "LSD: a Line Segment Detector," *IPOL*, vol. 2, pp. 35–55, 2012.

[32] R. Grompone von Gioi and G. Randall, "Unsupervised Smooth Contour Detection," *Image Processing On Line*, vol. 6, 2016.

[33] A. Gómez, G. Randall, and R. Grompone von Gioi, "A Contrario 3D Point Alignment Detection Algorithm," *Image Processing On Line*, vol. 7, 2017.

[34] X. Huang, W. Yang, H. Zhang, and G.-S. Xia, "Automatic ship detection in sar images using multi-scale heterogeneities and an a contrario decision," *Remote Sensing*, 2015.

[35] J. Lezama, G. Randall, J.-M. Morel, and R. Grompone von Gioi, "An Unsupervised Point Alignment Detection Algorithm," *Image Processing On Line*, vol. 5, 2015.

[36] J.-L. Lisani and S. Ramis, "A Contrario Detection of Faces with a Short Cascade of Classifiers," *Image Processing On Line*, vol. 9, 2019.

[37] J.-L. Lisani, S. Ramis, and F. J. Perales, "A contrario detection of faces: A case example," *SIAM Journal on Imaging Sciences*, vol. 10, 2017.

[38] Q. Bammey, "A contrario mosaic analysis for image forensics," in *Advanced Concepts for Intelligent Vision Systems, ACIVS 2023*. Springer, Aug. 2023.

[39] Q. Bammey, R. G. v. Gioi, and J.-M. Morel, "Reliable demosaicing detection for image forensics," in *27th European Signal Processing Conference (EUSIPCO)*, 2019, pp. 1–5.

[40] Q. Bammey, R. Grompone Von Gioi, and J.-M. Morel, "Demosaicing to detect demosaicing and image forgeries," in *2022 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2022, pp. 1–6.

[41] Y. Li, M. Gardella, Q. Bammey, T. Nikoukhah, J.-M. Morel, M. Colom, and R. Grompone von Gioi, "A contrario detection of h.264 video double compression," in *2023 IEEE International Conference on Image Processing (ICIP)*, 2023.

[42] Q. Bammey, R. Grompone von Gioi, and J.-M. Morel, "Automatic detection of demosaicing image artifacts and its use in tampering detection," in *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2018, pp. 424–429.

[43] T. Nikoukhah, J. Anger, M. Colom, J.-M. Morel, and R. Grompone von Gioi, "ZERO: a Local JPEG Grid Origin Detector Based on the Number of DCT Zeros and its Applications in Image Forensics," *IPOL*, 2021.

[44] T. Ehret, "Robust copy-move forgery detection by false alarms control," *arXiv preprint arXiv:1906.00649*, 2019.

[45] Q. Bammey, "Analysis and Experimentation on the ManTraNet Image Forgery Detector," *Image Processing On Line*, vol. 12, pp. 457–468, 2022, https://doi.org/10.5201/ipol.2022.431.

[46] H. Mareen, D. V. Bussche, F. Guillaro, D. Cozzolino, G. Van Wallendael, P. Lambert, and L. Verdoliva, "Comprint: Image forgery detection and localization using compression fingerprints," 2022. [Online]. Available: https://arxiv.org/abs/2210.02227

[47] D. Cozzolino, G. Poggi, and L. Verdoliva, "Splicebuster: A new blind image splicing detector," in *WIFS*, 2015.

[48] Q. Bammey, T. Nikoukhah, M. Gardella, R. G. von Gioi, M. Colom, and J.-M. Morel, "Non-semantic evaluation of image forensics tools: Methodology and database," in *WACV*, January 2022, pp. 3751–3760.

[49] D.-T. Dang-Nguyen, C. Pasquini, V. Conotter, and G. Boato, "Raise: A raw images dataset for digital image forensics," in *Proceedings of the 6th ACM Multimedia Systems Conference*, 2015, pp. 219–224.

[50] B. K. T. Ho, V. Y. Tseng, M. Ma, and D. T. Chen, "Mathematical model to quantify JPEG block artifacts," in *Medical Imaging 1993: Image Capture, Formatting, and Display*, Y. Kim, Ed., vol. 1897, International Society for Optics and Photonics. SPIE, 1993, pp. 269 – 275.

[51] B. Li, T.-T. Ng, X. Li, S. Tan, and J. Huang, "Revealing the trace of high-quality jpeg compression through quantization noise analysis," *IEEE Transactions on Information Forensics and Security*, vol. 10, no. 3, pp. 558–573, 2015.

[52] ——, "Statistical model of jpeg noises and its application in quantization step estimation," *IEEE Transactions on Image Processing*, vol. 24, no. 5, pp. 1471–1484, 2015.

[53] M. Gardella, T. Nikoukhah, Y. Li, and Q. Bammey, "The impact of jpeg compression on prior image noise," in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 2689–2693.

[54] N. Le and F. Retraint, "An improved algorithm for digital image authentication and forgery localization using demosaicing artifacts," *IEEE Access*, vol. 7, pp. 125 038–125 053, 2019.

[55] Q. Bammey, R. Grompone von Gioi, and J.-M. Morel, "Image Forgeries Detection through Mosaic Analysis: the Intermediate Values Algorithm," *Image Processing On Line*, vol. 11, pp. 317–343, 2021.

[56] S. Mandelli, N. Bonettini, P. Bestagini, and S. Tubaro, "Training cnns in presence of jpeg compression: Multimedia forensics vs computer vision," in *2020 IEEE International Workshop on Information Forensics and Security (WIFS)*. IEEE, 2020, pp. 1–6.