



HAL
open science

A Contrario Detection of H.264 Video Double Compression

Yanhao Li, Marina Gardella, Quentin Bammey, Tina Nikoukhah, Jean-Michel Morel, Miguel Colom, Rafael Grompone von Gioi

► **To cite this version:**

Yanhao Li, Marina Gardella, Quentin Bammey, Tina Nikoukhah, Jean-Michel Morel, et al.. A Contrario Detection of H.264 Video Double Compression. 2023 IEEE International Conference on Image Processing (ICIP), Oct 2023, Kuala Lumpur, Malaysia. pp.1765-1769, 10.1109/ICIP49359.2023.10222775 . hal-04225423

HAL Id: hal-04225423

<https://hal.science/hal-04225423>

Submitted on 2 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A CONTRARIO DETECTION OF H.264 VIDEO DOUBLE COMPRESSION

Yanhao Li^{*} Marina Gardella^{*} Quentin Bammey^{*} Tina Nikoukhah^{*}
Jean-Michel Morel[†] Miguel Colom^{*} Rafael Grompone von Gioi^{*}

^{*}Centre Borelli, ENS Paris-Saclay, Université Paris-Saclay, CNRS, France

[†]City University of Hong Kong, Department of Mathematics, Kowloon, Hong Kong

ABSTRACT

Video manipulation detection plays a vital role in modern multimedia forensics. In particular, double compression detection provides significant clues leading to the video edition history and hinting at potential malevolent manipulation. While such an analysis is well-understood on images, the research on this subject remains lacking in videos and existing methods are not yet able to reliably detect double-compressed videos. This work presents a novel method for identifying double compression in H.264 codec videos. Our technique exploits the periodicity of frame residuals caused by fixed Group of Pictures in the initial compression, and employs an *contrario* framework to minimize and control false detections. The proposed method can reliably detect double compression in videos. It does not require threshold tuning, thus enabling automatic detection. The code is available at https://github.com/li-yanhao/gop_detection.

Index Terms— Video double compression, video forensics, group of pictures, a contrario, deepfake detection

1. INTRODUCTION

The development of video post-processing software has made video edition a widespread practice. Though video editing can aim simply at enhancing the aesthetics, it can also be used for malicious purposes. Assessing the authenticity and integrity of videos has become an important task in modern societies. Video forensics research field emerged to address these concerns [1], [2]. Amongst passive forensic techniques, the detection of double compression can provide significant clues to recover the editing history of a video. Indeed, to manipulate a video, one must first decompress it, then perform the desired editions and finally recompress it. Recompression artifacts, though imperceptible to the human eye, can be detected by analyzing both the spatial statistics as is done for static images [3]–[6], and the temporal statistics hidden in the Group of Pictures (GOP) [7]–[17].

The GOP structure, which defines different types of frames and their order in a video, plays a crucial role in video encoding [18], and can be exploited in forensics [2], [19]. Frames called *intra pictures* (I-frames) are encoded independently of the other frames; the *predicted pictures* (P-frames) only encode changes relative to the previous frame; finally, *bidirectional predicted pictures* (B-frames) encode changes relative to its previous and subsequent frames. These frames have different properties: I-frames are the least compressible and independent of their neighbors, while B- and P-frames are more compressible and also dependent on neighboring frames. Analyzing the GOP structure of a video can thus help spot traces of a previous compression of a video.

We propose a method to detect whether a video has been recompressed by analyzing the potential abnormal artifacts left by the fixed GOP used during the first compression. Note that the fixed GOP size is used in H.264 baseline profile, and is also a common case when no scene changes occur [9] e.g. the background of deepfake faceswap videos.

2. RELATED WORKS

Wang and Farid [10] expose the temporal periodic increase of motion error with Discrete Fourier Transform to detect frame insertion or deletion in a recompressed video with same GOPs in both compressions. Stamm et al. [11] improve this method by adding the case of variable GOP in the second compression. Jiang et al. [12] use Markov statistics to find double quantization artifacts in MPEG-4 videos. Abbasi et al. [13] combine the DCT coefficients of I-frames with an SVM. In [14], the error between the true value of a pixel and the one estimated using all the other frames in a GOP is compared to a threshold to detect recompressed frames.

Several works also focus periodicity analysis to detecting video double compression. Vázquez-Padín et al. [15] analyze the temporal variation of intra-coded macroblocks and skipped macroblocks in P-frames, assuming that a fixed GOP was used in the first compression. Chen et al. [16] incorporate Prediction Residual Distribution features and processed periodicity analysis to detect double compression. Yao et al. [17] study the periodic features of the string of data bits and the S-MB to detect double encoding. He et al. [8] analyze the pe-

This work was supported by grants from ANR (APATE, ANR-22-CE39-0016), Horizon Europe VERA.AI (No. 101070093), Région Île-de-France and UDOPIA (ANR-20-THIA-0013). Centre Borelli is also a member of Université Paris Cité, SSA and INSERM.

riodic frame residuals in the background regions segmented by a motion vector field. Yao et al. [9] propose a strategy to detect recompressed videos with adaptive GOP by revealing the artifacts in the frame byte count sequence. Bestagini et al. [20] estimate the GOP using the similarity between a suspicious video and its additionally compressed version.

Other works focus on HEVC codec, including a deep-learning based approach [21], video degradation analysis [22], statistical analysis of Prediction Units (PU) [23] and an SVM classifier on PU and DCT coefficient features [24].

3. ANALYSIS OF DOUBLE COMPRESSION

We assume a constant GOP size during compression, and start by studying the GOP structure based on I and P frames only (see Fig. 1) as considered in [9], [15]–[17]. Each GOP starts with an I-frame followed by P-frames. A singly encoded P-frame only encodes the difference w. r. t. its reference frame, also singly encoded, as prediction residual, then stores its quantized version. The lossy quantization of frame residuals correlates a P-frame with its previous frame.

We also assume the GOP size of the second compression is different from the first one. After the second compression, there are two kinds of P-frames: I-P frames and P-P frames. An I-P (resp. P-P) frame is an I-frame (resp. P-frame) in the first compression relocated as a P-frame in the second compression. Since an I-P frame is not correlated to its reference frame *before* the second compression while a P-P frame is, the residuals *after* the second compression in I-P frames tend to be much larger than in P-P frames. This phenomenon is depicted in Fig. 1, where abnormal periodic peaks appear after the second compression. In addition, given that the GOP size is constant and the I-frames are equally spaced in the first compression, the I-P frames in the second compression are also equally spaced and form periodic residual peaks with period equal to the primary GOP size. We can thus detect double compression by finding a sequences of periodic residual peaks in the P-frames.

4. PROPOSED METHOD

Our *a contrario* method consists in detecting periodic residual peaks in P-frames of a video to decide whether the video is recompressed. Let $R_t \in \mathbb{R}^{H \times W}$ be the prediction residual of a P-frame at time t ; we use $r_t = \frac{1}{HW} \sum_{i=0}^{H-1} \sum_{j=0}^{W-1} |R_t(i, j)|$ as the frame residual. Different from previous works [11][16], we use the Cr color space instead of the luminance space because motion residuals in Cr plane are downsampled and the residual peaks of I-P frames are slightly more distinctive.

4.1. A *contrario* detection framework

The *a contrario* framework [25]–[27] is based on the *non-accidentalness principle* which states that a structure is relevant whenever a large deviation from randomness occurs. The main idea is to control the Number of False Alarms (NFA) of an event under the null model H_0 . Let

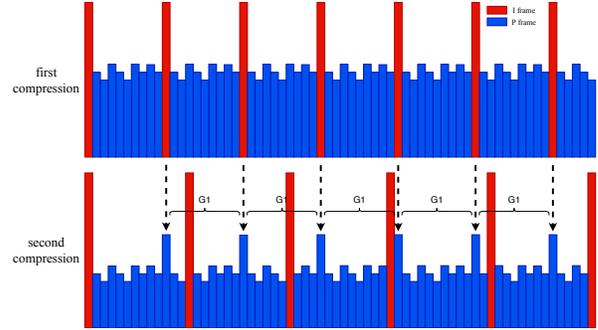


Fig. 1. When a video is compressed, *intra* frames (blue) are encoded independently of other frames, while *predicted* frames (red) only encode changes relative to previous frames. If the image is compressed twice, the prediction residuals of predicted frames are higher for the frames that were *intra* frames in the first compression than for those that were already predicted frames. This abnormal periodic peak can be detected as evidence of the double compression.

$\mathbf{r} := \{r_t : t\text{-th frame is not I-frame}\}$ be a *multivariate random variable* representing a sequence of prediction residuals of P-frames. Our null hypothesis H_0 is that the video was not recompressed, so that there are no periodic residual peaks in P-frames; P-frames close to each other should have similar residuals. If there was a double compression, there should be a periodic sequence of P-frame residuals $S(p, b, \mathbf{r}) \triangleq \{r_b, r_{b+p}, r_{b+2p}, \dots\} \cap \mathbf{r}$ starting at r_b with period p that are more likely to have larger values than their neighbors. Given a candidate sequence $S(p, b, \mathbf{r})$, we count the number of its elements which are greater than their d neighboring P-frame residuals on each side and note it as

$$k(p, b, \mathbf{r}) \triangleq \sum_{r \in S(p, b, \mathbf{r})} \mathbb{1}_{\{r \geq \max B_d(r)\}}, \quad (1)$$

where $B_d(r) \subset \mathbf{r}$ is the set of residuals at a neighborhood of r considering d P-frames before and d P-frames after in \mathbf{r} .

If there are sufficient peak elements in the sequence $S(p, b, \mathbf{r})$, we say the sequence is a salient observation. Thus we need a threshold $\kappa(p, b)$ to validate the sequence if $k(p, b, \mathbf{r}) \geq \kappa(p, b)$. Note that $\kappa(p, b)$ depends on each sequence parameters (p, b) . Let C be the candidate set containing all possible pairs of (p, b) , then the total number of detections is given by:

$$D(\mathbf{r}) \triangleq \sum_{(p, b) \in C} \mathbb{1}_{\{k(p, b, \mathbf{r}) \geq \kappa(p, b)\}}. \quad (2)$$

The *a contrario* framework requires κ to be set in such a way that the expected number of detections $D(\mathbf{r})$ under H_0 is smaller than a threshold ε :

$$\mathbb{E}[D(\mathbf{r})] = \sum_{(p, b) \in C} \mathbb{P}\left(k(p, b, \mathbf{r}) \geq \kappa(p, b)\right) < \varepsilon. \quad (3)$$

There are many choices of κ to achieve this adjustment. Here the Bonferroni correction [28] is applied. We group all the possible $(p, b) \in C$ by period and divide ε in equal parts

among all them, then in equal parts among all the possible offsets for each period, which, for a given p , are $b \in [0, p - 1]$. Since we need d neighbor residuals on both sides of a tested residual, the distance between two tested residuals should not be smaller than $2d + 1$. The maximum period should be smaller than the half of sequence length. Given a video of n frames, the possible periods are thus $p \in [2d + 1, \lfloor \frac{n-1}{2} \rfloor]$. With this partition, a candidate (p, b) is assigned with $\kappa(p, b)$:

$$\kappa(p, b) = \min \left\{ \eta, \mathbb{P} \left(k(p, b, \mathbf{r}) \geq \eta \right) < \frac{\varepsilon}{N(p, b)} \right\}, \quad (4)$$

where $N(p, b)$ is the number of possible periods times the number of possible offsets given p , namely $N(p, b) = (\lfloor \frac{n-1}{2} \rfloor - 2d)p$. By doing so, Eq. 3 is satisfied.

Consider now an observed sequence of P-frame residuals X as a realization of \mathbf{r} . A candidate (p, b) gives us the number of peaks $K = k(p, b, X)$. Instead of directly comparing $\kappa(p, b)$ and K , an equivalent way is to compute the NFA of this candidate and validate it if $\text{NFA}(p, b, X) < \varepsilon$:

$$\text{NFA}(p, b, X) \triangleq N(p, b) \cdot \mathbb{P}(k(p, b, \mathbf{r}) \geq K). \quad (5)$$

It still remains to compute the term $\mathbb{P}(k(p, b, \mathbf{r}) \geq K)$. Under the null hypothesis H_0 , the probability that a particular residual has the largest value among n residuals is $1/n$. Thus, a tested residual is observed to be a peak larger than all of its $2d$ neighbors with probability

$$\rho = \mathbb{P} \left(r \geq \max B_d(r) \right) = \frac{1}{2d + 1}. \quad (6)$$

Since any two tested residuals use disjoint neighborhoods for peak validation, each observation of a peak residual can be considered independent. Then, the number of peak elements in a sequence $S(p, b, X)$ follows a binomial distribution $k(p, b, \mathbf{r}) \sim \mathcal{B}(\#S(p, b, \mathbf{r}), \frac{1}{2d+1})$ and the NFA is given by:

$$\begin{aligned} \text{NFA}(p, b, X) &= N(p, b) \cdot \mathbb{P}(k(p, b, \mathbf{r}) \geq K) \\ &= \left(\left\lfloor \frac{n-1}{2} \right\rfloor - 2d \right) \cdot p \cdot \mathbf{B} \left(K, \#S(p, b, \mathbf{r}), \frac{1}{2d+1} \right) \end{aligned} \quad (7)$$

where n is the number of frames of the video, d is the range of each test neighborhood, $\#S(p, b, \mathbf{r})$ is the length of the tested periodic sequence, K is the observed number of residual peaks, and \mathbf{B} is the tail of the binomial law: $\mathbf{B}(l, m, \rho) = \sum_{i=l}^m \binom{m}{i} \rho^i (1 - \rho)^{m-i}$.

4.2. Adaptation to videos with B-frames

The initial method is effective for videos containing only I and P frames, where abnormal residual increases occur strictly in periodic I-P frames. However, for videos incorporating B-frames, an I-frame might be recast as a B-frame during the second compression, resulting in periodic residual increases in both I-P and I-B frames. Direct residual comparisons between P-frames and B-frames are impractical due to variable compression rates. Nonetheless, any abnormal residual increase in an I-B frame will also emerge in the next P-frame (refer to Fig. 2), as the second compression causes the follow-

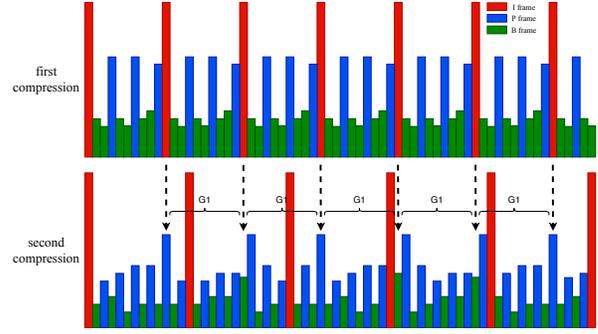


Fig. 2. A schematic diagram of artifacts in the prediction residuals of P-frames in a video compressed with B-frames, presented by quasi-periodic abnormal peaks. At the 2nd, 4th and 5th arrow the abnormal residual increase occurs in the subsequent P-frame of an I-B frame.

ing P-frame to reference a frame from a different GOP during the initial compression, thus enhancing prediction residuals.

To adapt our method to videos with B frames, we replace the direct testing of candidate sequences $S(p, b, \mathbf{r}) = \{r_b, r_{b+p}, r_{b+2p}, \dots\} \cap \mathbf{r}$ with an indirect evaluation of $\tilde{S}(p, b, \mathbf{r}) \triangleq \{\tilde{r}_b, \tilde{r}_{b+p}, \tilde{r}_{b+2p}, \dots\} \cap \mathbf{r}$, where $r_i \mapsto \tilde{r}_i$ assigns each residual to its immediate subsequent P-frame residual. \tilde{r}_i is discarded if no subsequent P-frame of r_i is found or if there is an I-frame between r_i and \tilde{r}_i covering the abnormal residual increase.

5. EXPERIMENTS

To compare the proposed method with [15]–[17], we first selected 19 uncompressed YUV sequences¹. Each video was clipped to no more than 400 frames. Following the settings in [16] and [17] and using *ffmpeg* software with *libx264* encoder, the first compression was processed with different constant bitrates $B1 \in \{300, 700, 1100\}$ kbps and GOP sizes $G1 \in \{10, 15, 30, 40\}$, while for the second compression the bitrates were $B2 \in \{300, 700, 1100\}$ kbps and the GOP sizes were $G2 \in \{9, 16, 33, 50\}$. Note that the higher the bitrate, the lower the compression. In total, the constructed dataset has 228 singly compressed videos and 2736 doubly compressed videos. Considering that the compared methods only work on videos without B frames, the videos were first compressed only with I and P frames for comparison. Besides, the compared methods only detect recompressed videos without any time shift and only look for periodic signals starting at the first frame. Therefore, the method is set to only detect periodic sequences with offsets $b = 0$. The number of neighbors on each side for validating a peak residual is set to $d = 3$.

Since the compared detectors all rely on specific thresholds, we first compared their areas under the Receiver Operating Characteristic curves (AUROC), which delivers a threshold-independent comparison. We grouped the videos by different bitrates $B2$ and computed the AUROC for each

¹ *akiyo, bridge_close, bridge_far, city, crew, deadline, flower_garden, football(a), foreman, galleon, harbour, ice, highway, mad900, mthr_dotr, paris, students, soccer and sign_irene* from <https://media.xiph.org/video/derf/>

B2	300	700	1100	all
Chen	0.645	0.789	0.866	0.767
Yao	0.281	0.667	0.833	0.593
Vázquez-Padín	0.866	0.932	0.945	0.914
Proposed	0.930	0.967	0.967	0.955

Table 1. AUROCs of double compression detection for different bitrates B2 of the second compression.

subgroup. Tab. 1 shows our method outperforms the other ones, especially when the second compression is high.

In addition, we compute the Precision-Recall curves for all methods. To avoid imbalanced classification, the recompressed videos are partitioned into subsets of equal numbers of singly compressed videos, each subset is merged to the same set of single compressed videos, the average score is computed. Fig. 3 shows the curves obtained for each method. The proposed method achieves the best area under the PR curves, as seen in Tab 3. Our method is also best suited for forensic applications, since it enables a controlled and low false positive rate while keeping the best recall.

We further investigated the comparison with threshold tuning. The videos originated from 12 of the raw sequences (1872 videos in total) were selected as training test to find the empirical parameters of [15]–[17], and the videos from the other 7 sequences (1092 videos in total) were used as test set. Each compared method was assigned with a threshold such that the precision for the training set is 95%, then performed double compression detection with the same threshold on the test set. As for our method, the threshold ϵ represents an upper bound of the expected number of false detections. We would normally require a mean number of false detections smaller than 1. Due to the discrete nature of the binomial law, the average number of false detections is actually much smaller than the upper bound ϵ [29]. Therefore, we tested the method with ϵ set to 1 and 0.1. The precision, recall and F1 score, in Tab. 2, show that, with a simple and reasonable choice of ϵ , our method outperforms the other ones for all the metrics. This is because the used meaningfulness threshold ϵ does not depend on the characteristics of the videos, whereas the compared methods require empirical parameters adapted to each video. Thus, these methods have less generalizability.

Method	Precision	Recall	F1
Chen [16]	0.976	0.578	0.726
Yao [17]	0.660	0.419	0.512
Vázquez-Padín [15]	0.667	0.075	0.135
Proposed, $\epsilon = 1$	1.000	0.969	0.984
Proposed, $\epsilon = 0.1$	1.000	0.958	0.979

Table 2. Precisions, recalls and F1 scores on the test set. The compared methods use thresholds pre-selected in the training set, while our method only uses manually chosen values of ϵ as the upper bound of the expected number of false detections.

Method	Chen	Yao	Vázquez-Padín	Proposed
AUPRC	0.809	0.651	0.897	0.937

Table 3. Areas under the Precision-Recall curves, related to the solid curves shown in Fig. 3.

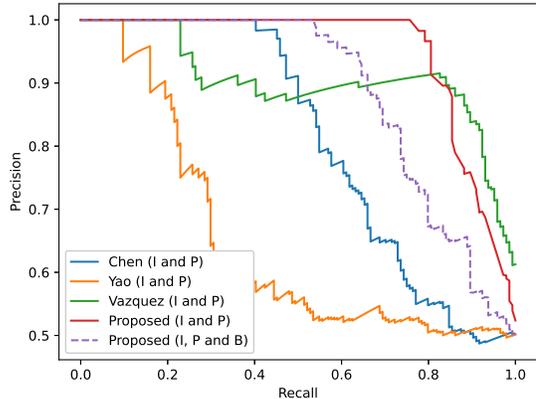


Fig. 3. Precision-Recall curves including all the videos of different encoding bitrates and GOPs. Solid curves correspond to the videos encoded with only I and P frames while the dashed curve is associated to the same videos encoded with I, P and B frames for our method.

Furthermore, the change of ϵ from 1 to 0.1 only has a minor impact on the performance of our method, indicating its low dependency on this threshold.

Finally, to evaluate the performance of the proposed method on videos containing B-frames, the same sequences were compressed with the same settings except for the use of B-frames. Then detection was performed using the adaptation described in Sec. 4.2. The PR curve obtained is shown by the dashed curve in Fig. 3. As can be seen the presence of B-frames increases the difficulty of the double compression detection, but at 100% precision our method is still able to retrieve more than 50% of the recompressed videos.

6. CONCLUSION

We proposed to detect video double compression in H.264 codec by detecting periodicity of frame residuals of P-frames, based on the *a contrario* statistical theory. The proposed method relies on the assumption that the GOP of the first and second compression is different, and our simple frame residual is susceptible to strong motions. We further adapted the method to videos with B-frames. Our experiments show this method beats the SOTA methods without threshold tuning. In the future, the proposed detector could be completed with learning-based methods such as positional learning [30]–[32] and combined with other distinctive frame-wise features (e.g. [15]) to further improve its performance. It could be extended to other video codecs and be adapted towards video forgery detection especially for deepfake videos.

References

- [1] K. Sitara and B. M. Mehtre, "Digital video tampering detection: An overview of passive techniques," *Digital Investigation*, vol. 18, pp. 8–22, 2016.
- [2] N. Shelke and S. Singh Kasana, "A comprehensive survey on passive techniques for digital video forgery detection," *Multimedia Tools and Applications*, 2021.
- [3] T. Nikoukhah, J. Anger, T. Ehret, M. Colom, J.-M. Morel, and R. G. von Gioi, "Jpeg grid detection based on the number of dct zeros and its application to automatic and localized forgery detection," in *IEEE/CVF CVPRW*, 2019.
- [4] Q. Bammey, "Jade owl: Jpeg 2000 forensics by wavelet offset consistency analysis," in *IEEE ICIVC*, 2023.
- [5] M. Gardella, T. Nikoukhah, Y. Li, and Q. Bammey, "The impact of jpeg compression on prior image noise," in *ICASSP*, 2022.
- [6] Q. Wang and R. Zhang, "Double jpeg compression forensics based on a convolutional neural network," *EURASIP JIS*, 2016.
- [7] W. Wang and H. Farid, "Exposing digital forgeries in video by detecting double mpeg compression," in *MM&Sec*, ACM, 2006.
- [8] P. He, X. Jiang, T. Sun, and S. Wang, "Double compression detection based on local motion vector field analysis in static-background videos," *JVCIR*, 2016.
- [9] H. Yao, R. Ni, and Y. Zhao, "Double compression detection for h. 264 videos with adaptive gop structure," *Multimedia Tools and Applications*, 2020.
- [10] W. Wang and H. Farid, "Exposing digital forgeries in video by detecting double quantization," in *ACM MM&Sec'09*, 2009.
- [11] M. C. Stamm, W. S. Lin, and K. R. Liu, "Temporal forensics and anti-forensics for motion compensated video," *IEEE TIFS*, 2012.
- [12] X. Jiang, W. Wang, T. Sun, Y. Q. Shi, and S. Wang, "Detection of double compression in mpeg-4 videos based on markov statistics," *IEEE SPL*, 2013.
- [13] J. A. Aghamaleki and A. Behrad, "Malicious inter-frame video tampering detection in mpeg videos using time and spatial domain analysis of quantization effects," *Multimedia Tools and Applications*, 2017.
- [14] A. V. Subramanyam and S. Emmanuel, "Pixel estimation based video forgery detection," in *ICASSP*, IEEE, 2013.
- [15] D. Vazquez-Padin, M. Fontani, T. Bianchi, P. Comesaña, A. Piva, and M. Barni, "Detection of video double encoding with gop size estimation," in *WIFS*, IEEE, 2012.
- [16] S. Chen, T. Sun, X. Jiang, P. He, S. Wang, and Y. Q. Shi, "Detecting double h. 264 compression based on analyzing prediction residual distribution," in *IWDW 2016*, Springer, 2017.
- [17] H. Yao, S. Song, C. Qin, Z. Tang, and X. Liu, "Detection of double-compressed h. 264/avc video incorporating the features of the string of data bits and skip macroblocks," *Symmetry*, 2017.
- [18] I. E. Richardson, *The H. 264 advanced video compression standard*. John Wiley & Sons, 2011.
- [19] Y. Li, M. Gardella, Q. Bammey, *et al.*, "Video signal-dependent noise estimation via inter-frame prediction," in *2022 IEEE International Conference on Image Processing (ICIP)*, 2022.
- [20] P. Bestagini, S. Milani, M. Tagliasacchi, and S. Tubaro, "Codec and gop identification in double compressed videos," *IEEE TIP*, 2016.
- [21] K. Uddin, Y. Yang, and B. T. Oh, "Deep learning based hevc double compression detection," 2020.
- [22] X. Jiang, Q. Xu, T. Sun, B. Li, and P. He, "Detection of hevc double compression with the same coding parameters based on analysis of intra coding quality degradation process," *IEEE TIFS*, 2020.
- [23] X. Jiang, P. He, T. Sun, and R. Wang, "Detection of double compressed hevc videos using gop-based pu type statistics," *IEEE Access*, 2019.
- [24] Z. Li, R.-s. Jia, Z. Zhang, X. Liang, and J. Wang, "Double hevc compression detection with different bitrates based on co-occurrence matrix of pu types and dct coefficients," 2017.
- [25] A. Desolneux, L. Moisan, and J.-M. Morel, *From gestalt theory to image analysis: a probabilistic approach*. Springer Science & Business Media, 2007.
- [26] J. Wagemans, "Perceptual use of nonaccidental properties," *Canadian Journal of Psychology*, 1992.
- [27] M. K. Albert and D. D. Hoffman, "Genericity in spatial vision.," 1995.
- [28] Y. Hochberg and A. Tamhane, *Multiple Comparison Procedures*. John Wiley, 1987.
- [29] R. Grompone von Gioi and J. Jakubowicz, "On computational gestalt detection thresholds," *Journal of Physiology-Paris*, 2009.
- [30] 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.
- [31] Q. Bammey, R. G. von Gioi, and J.-M. Morel, "Forgery detection by internal positional learning of demosaicing traces," in *IEEE/CVF WACV*, 2022.
- [32] Q. Bammey, "A contrario mosaic analysis for image forensics," in *ACIVS*, Springer, 2023.