AN INFORMATION THEORETIC FRAMEWORK FOR ORDER OF OPERATIONS FORENSICS

Xiaoyu Chu[*], *Yan Chen*[*#] *and K. J. Ray Liu*[*]

[*]Department of Electrical and Computer Engineering, University of Maryland, College Park
[#]School of Electronic Engineering, University of Electronic Science and Technology of China

ABSTRACT

To verify the authenticity of easily manipulated multimedia content, forensic researchers have proposed many techniques to estimate the processing history of given multimedia content. When multiple operations may be applied on multimedia content, a complete processing history would involve the information of not only what manipulation operations have been applied, but also in what order they were applied. However, there are few works considering the problem of detecting the order of operations. Moreover, due to the interplay among operations, the order of operations may not always be detectable. This leads to a fundamental question of when we can or cannot detect the order of operations. In this paper, we propose an information theoretical framework to answer this question. Specifically, we formulate the problem of detecting the order of operations into a multiple hypotheses testing problem. Then, we propose an information theoretical framework to characterize the relationship between the true hypothesis and the detected hypothesis. Under this framework, we propose a mutual information based criterion to determine the detectability of the order of operations. Furthermore, conditional fingerprints are defined in this framework to understand why the order of operations is not always detectable. The detection of the order of resizing and blurring is examined in this paper, where the order detection scheme has been proposed and the effectiveness of our framework has been demonstrated by simulations.

Index Terms— Digital Forensics, Order of Operations, Information Theory, Resizing and Blurring

1. INTRODUCTION

With the help of various multimedia editing tools, manipulating multimedia content has become very easy. In order to verify the authenticity of multimedia content, forensic researchers have developed many techniques to estimate the processing history of the given multimedia content. Specifically, we are able to identify the use of different manipulation operations such as compression [1–3], resizing [4, 5], contrast enhancement [6], blurring [7–9] and so on [10–12].

Most of these forensic techniques, however, made assumptions that only one operation may be applied on the multimedia content. For scenarios where multiple operations are considered, there have been some techniques proposed to detect the existence of specific operations in an operation chain. For example, in [13] and [14], forensic techniques have been proposed to detect double compression in the operation chain of two compressions with resizing or contrast enhancement in between, respectively. In [15], an improved contrast enhancement detector was proposed to detect this operation when it was applied on previously JPEG compressed images. Furthermore, authors in [16] were able to recover the compression history when full-frame linear filtering is applied after JPEG.

While it is important to detect the existence of each operation in an operation chain, detecting the order of how these operations have been applied is even more crucial to obtain the complete processing history of multimedia content. Furthermore, the knowledge of the order of operations can help us investigate when the multimedia content was manipulated and who manipulated it.

Yet, few works have been done on detecting the order of operations. Authors in [17] have examined the detection of the order of resizing and contrast enhancement and proposed a forensic technique to detect their order. However, due to the effect of later applied operation on the fingerprints of earlier applied operations, the order of operations in an operation chain is not always detectable. For example, if JPEG compression or Gaussian noise is applied after contrast enhancement, the fingerprints of contrast enhancement would be too weak to be detected [15].

Therefore, a natural question would be "when can we or cannot we detect the order of operations?" Authors in [18] proposed a measure of distinguishability on simple hypotheses problems. While in this paper, we answer the question by formulating the problem of detecting the order of operations into a more general multiple hypotheses testing problem. Then, we propose an information theoretical framework and mutual information based criteria to determine the condition of when we can distinguish all considered hypotheses. The detection of the order of resizing and blurring is examined in this paper to demonstrate the effectiveness of our framework and criteria.

2. DETECTABILITY OF THE ORDER OF OPERATIONS

In [17], a forensic technique has been proposed to detect the order of resizing and contrast enhancement, and simulation results have shown that their order can be successfully detected. However, in general, the order of operations may not always be detectable. Let us consider an example of detecting the order of resizing and blurring. In this problem, we consider the following five hypotheses for possible processing history of a given image.

$$\begin{aligned} &H_0 : \text{It is unaltered,} \\ &H_1 : \text{It is altered by } A \text{ only,} \\ &H_2 : \text{It is altered by } B \text{ only,} \\ &H_3 : \text{It is altered by } B \text{ then } A, \\ &H_4 : \text{It is altered by } A \text{ then } B, \end{aligned} \tag{1}$$

where A and B denote the operations of resizing and blurring respectively.

We have found that the fingerprints of each hypothesis on the discrete Fourier transform (DFT) of an image's p-map [4] can be used to contrast different hypotheses. As they are shown in Fig. 1(b)

This work is supported in part by the NSF grant CCF1320803
Email: {cxygrace,yan,kjrliu}@umd.edu.

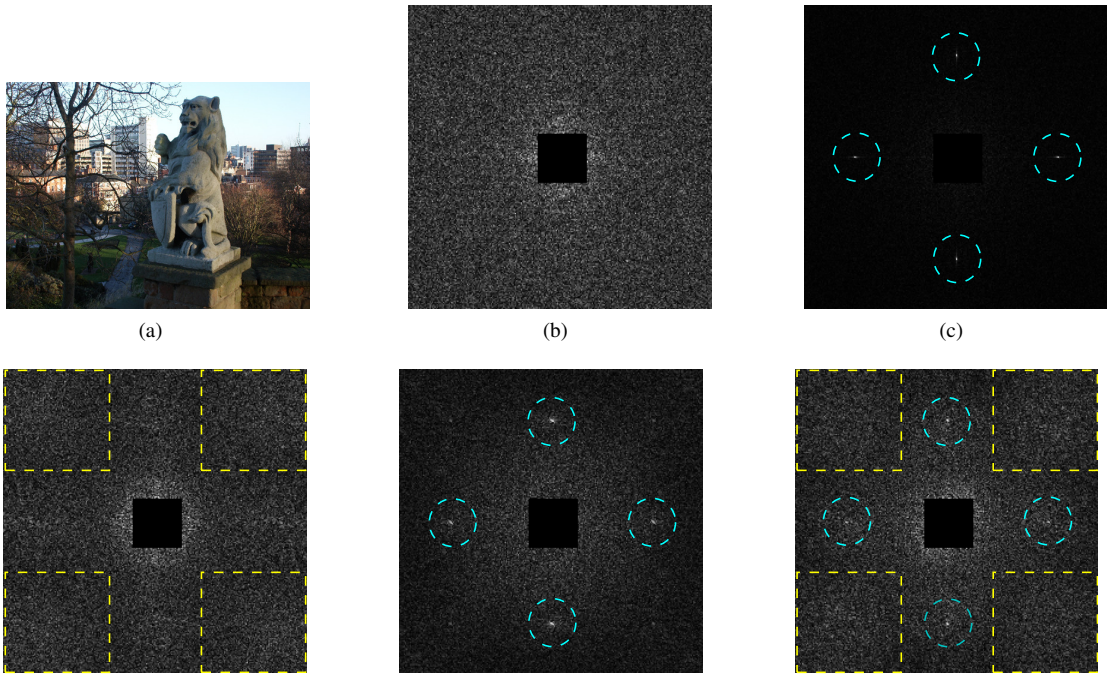

(a) (b) (c) (d) (e) (f)

Fig. 1: Fingerprints for detecting the order of resizing and blurring. (a) and (b) are the original image and the DFT of its p-map, respectively. (c) - (f) show the DFT of the p-map of (c) the resized image, (d) the blurred image, (e) the blurred then resized image, and (f) the resized then blurred image. Resizing factor is 1.5 (upscaling). Gaussian blur is used with variance 1. Regions of interests are highlighted by dotted squares and circles.

- 1(f), if resizing is applied on the image, four distinct peaks can be revealed from the DFT of an image's p-map. If blurring is applied as the last operation, an increase of noise energy can be observed in high frequency component of the DFT of an image's p-map. Then, if resizing is applied before blurring, both fingerprints of resizing and blurring will be seen, while for the case where resizing is applied after blurring, fingerprints of blurring will be hardly detectable. For both hypotheses, the DFTs of the p-map are more noisy than that of a resized only image.

Based on these fingerprints, one may design forensic techniques to contrast different hypotheses in (1). However, there are some cases where the fingerprints are too weak to be detected. Fig. 2 shows a confusing example where blurring effect is weaker than that in Fig. 1. We can see that the fingerprints of resizing then blurring and blurring then resizing are hardly distinguishable. In this case, we may not be able to detect the order of resizing and blurring.

3. INFORMATION THEORETICAL FRAMEWORK AND MUTUAL INFORMATION BASED CRITERIA

In order to determine when we can or cannot detect the order of operations, we formulate the order detection problems into multiple hypotheses testing problems. Given the set of hypotheses the given multimedia content may belong to, forensic investigators first find fingerprints that can be used to contrast these hypotheses. Then, based on these fingerprints, features can be extracted from the multimedia content, and finally detectors will be used to obtain the detected hypothesis. The process has been shown in Fig. 3.

Let $\mathcal{H} = \{H_0, H_1, ..., H_{M-1}\}$ denote the set of considered hypotheses in a multiple hypotheses testing problem. Then the true hypothesis and the detected hypothesis, denoted as H and $\hat{H}$ respectively, belong to this set. Based on certain features, detectors with tunable parameters $\underline{\theta}$, denoted as $d_{\underline{\theta}}$, can be developed to contrast the different hypotheses. For each choice of $\underline{\theta}$, we represent the performance of the specific detector using a transition probability matrix $\mathbf{T}(\underline{\theta})$ with each element denoting the conditional probability of a detected hypothesis given a true hypothesis, i.e.,

$$\mathbf{T}_{i,j}(\underline{\theta}) = \mathbb{P}_{\underline{\theta}}(\hat{H} = H_j | H = H_i), \quad 0 \leq i, j < M. \tag{2}$$

With this representation, we have proposed a feature dependent abstract channel between the true hypothesis and the detected hypothesis. The channel characterization is specified by the parameters of the set of detectors, as it is shown in Fig. 3. Note that comparing with our earlier work [19] which examined more fundamental relationships between true hypotheses and features and proposed the concept of forensicability, this work examines the relationship between true hypotheses and estimated hypotheses and gives a more tractable way to determine the detectability of order of operations.

Given that different parameters of the detector yield different detection performance, in order to determine whether we can distinguish hypotheses in (1), we first find the best detector and then check if the best detector can distinguish all considered hypotheses. Under the above information theoretical framework, we propose a mutual information based criterion to obtain the best performed detector whose detected hypotheses contain the maximum information about the true hypotheses.

Definition 1 *For detectors $d_{\underline{\theta}_1}$ and $d_{\underline{\theta}_2}$, whose transition probability matrices are $\mathbf{T}(\underline{\theta}_1)$ and $\mathbf{T}(\underline{\theta}_2)$ respectively, assume uniform priors of the true hypothesis H, then, detector $d_{\underline{\theta}_1}$ is better than $d_{\underline{\theta}_2}$, w.r.t. the mutual information criterion, when*

$$I_{\mathbf{T}(\underline{\theta}_1)}(H; \hat{H}) > I_{\mathbf{T}(\underline{\theta}_2)}(H; \hat{H}), \tag{3}$$

where $I(H; \hat{H})$ denotes the mutual information between H and $\hat{H}$.

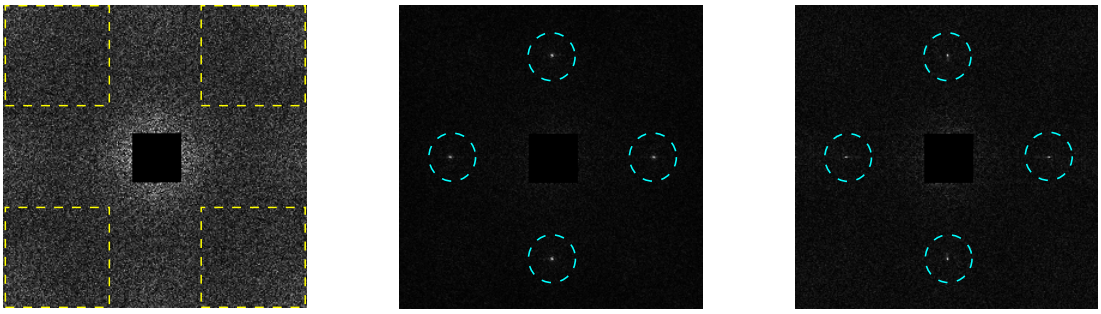

(a) (b) (c)

Fig. 2: A confusing example that we may not be able to detect the order. Plotted are DFTs of the p-map of (a) the blurred image, (b) the blurred then resized image, and (c) the resized then blurred image when resizing factor is 1.5 and the variance of Gaussian blur is 0.7.

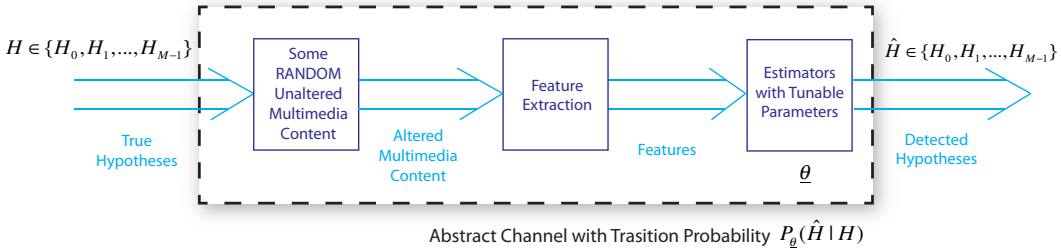

Fig. 3: A typical process of contrasting different hypotheses.

Then, we examine whether this best detector can successfully distinguish all hypotheses by checking if the detection rate is higher than any misdetection rate for each hypothesis.

Definition 2 *Under the mutual information criterion, all hypotheses can be distinguished by detectors $d_{\underline{\theta}}, \underline{\theta} \in \mathbb{R}^k$, if and only if*

$$H_i = \arg\max_{t \in \mathcal{H}} \mathbb{P}_{\underline{\theta}^*}(\hat{H} = t | H = H_i), \ \forall i = 0, 1, ..., M-1, \quad (4)$$

where $\underline{\theta}^*$ are parameters of the best detector w.r.t. the mutual information criterion. That is,

$$\underline{\theta}^* = \arg\max_{\underline{\theta}} I_{\mathbf{T}(\underline{\theta})}(H; \hat{H}). \quad (5)$$

If the best detector can distinguish all hypotheses, then considered hypotheses in (1) can be distinguished and the order of operations can be detected; If the best detector cannot, then no detector can distinguish the hypotheses and the order cannot be detected.

Furthermore, by checking which condition in (4) is violated, we can understand the reason of why the order cannot be detected, that is, which hypotheses are confused and cannot be distinguished. In this framework, fingerprints of single manipulation operations and conditional fingerprints in operation chains can be defined as well.

Definition 3 *Consider an operation chain and its corresponding hypothesis, denoted as $\underline{S}_i$ and H_i respectively. Let $\underline{S}_\emptyset$ and $H_\emptyset$ denote the empty operation chain, and the hypothesis of unaltered multimedia content. If $\underline{S}_i \neq \underline{S}_\emptyset$, then* the fingerprints of *$\underline{S}_i$ are a set of features that can be used to distinguish $\{H_i, H_\emptyset\}$. Next, we consider another operation chain, denoted as $\underline{S}_j$. If $\underline{S}_i$ is a sub-chain of $\underline{S}_j$, let $\underline{S}_{j\setminus i}$ denote the operation chain of $\underline{S}_j$ excluding $\underline{S}_i$. $H_{j\setminus i}$ is denoted as the corresponding hypothesis of $\underline{S}_{j\setminus i}$. Then, the conditional fingerprints of $\underline{S}_i$ given $\underline{S}_j$ are a set of features that can be used to distinguish the following hypotheses:*

$$\{H_{j\setminus i}, H_i, H_j\}.$$

To better understand the difference between fingerprints and conditional fingerprints, we give an example. Let $\underline{S}_i$ and $\underline{S}_j$ denote the operation chain of contrast enhancement only and contrast enhancement then resizing, respectively. Then, $\underline{S}_{j\setminus i}$ represents the operation chain of resizing only. The fingerprints of $\underline{S}_i$ can be the high frequency components of the DFT of the pixel histogram [6]. However, these cannot be the conditional fingerprints of $\underline{S}_i$ given $\underline{S}_j$ because $H_{j\setminus i}$ and H_j cannot be distinguished by these fingerprints. In [17], the conditional fingerprints of $\underline{S}_i$ given $\underline{S}_j$ contains two features. One is the maximum gradient of the periodogram of the Fourier transformed p-map, which is the fingerprint of resizing. The other feature is the distance of normalized histograms between the full image and the down-sampled image [17]. By using these two features, we can distinguish hypotheses $H_{j\setminus i}$, H_i, and H_j.

4. DETECT THE ORDER OF RESIZING AND BLURRING

In order to demonstrate the effectiveness of our proposed framework and criteria, we use the example of detecting the order of resizing and blurring. We first develop a forensic technique to distinguish the five hypotheses in (1) based on the fingerprints shown in Fig. 1.

Two features have been used in our detectors. One is the peak signal to noise ratio (PSNR) feature used to capture the four peaks in the DFT of an image's p-map and also the noise energy around these peaks. The other feature is to capture the noise energy pattern around the corners of the DFT of an image's p-map.

Let $Z = \{Z_{m,n}\}$ denote the magnitudes of the DFT of a p-map. The origin is located at the upper left corner of the matrix with size a by a. We first calculate the PSNR for the left peak in the DFT of an image's p-map. Let y_l denote the magnitudes of the left part of the horizontal line in Z, i.e., $y_l = [Z_{\lfloor a/2 \rfloor + 1, 1}, Z_{\lfloor a/2 \rfloor + 1, 2}, ..., Z_{\lfloor a/2 \rfloor + 1, \lfloor a/2 \rfloor + 1 - \alpha}]$, where α is the size of the mask window used in the center of Z to eliminate the effect from low frequency component.

Let x denote the index of the vector y_l. We have observed that the noise mean increases with x. To eliminate this nonuniform noise mean, we first use the following linear regression model to fit y_l,

$$y_l = a_1 x + b_1 + n. \tag{6}$$

After estimating the parameters $\hat{a}_1$ and $\hat{b}_1$, the PSNR feature will be calculated on the difference signal $d_l = y_l - \hat{a}_1 x - \hat{b}_1$,

$$\text{PSNR}_l = \frac{y_p}{\text{mean}_{0<|x-x_p|<\varepsilon}\Big(\big|d_l(x)\big|\Big)}, \tag{7}$$

where x_p and y_p are the index and magnitude of the peak respectively, and ε is a small positive constant.

Similar procedure can be applied on the other three peaks at right, top, and bottom. The final PSNR feature is the maximum value of the four PSNRs.

To obtain the noise energy pattern at the corners of Z, we first obtain the energy signal E as

$$E = Z \otimes \mathbf{1}_w, \tag{8}$$

where $\mathbf{1}_w$ is an all one matrix of size w by w, and $\otimes$ is a convolution operator. Then, the noise pattern signal is calculated as the mean signal of the boundary signals of E,

$$y_e(x) = \frac{\big(E_{v,a/2+x} + E_{a/2+x,a-v} + E_{a-v,a/2-x} + E_{a/2-x,v}\big)}{4}, \tag{9}$$

where $v - a/2 \le x < a/2 - v$ and $v = \lceil w/2 \rceil + 1$.

To characterize the increase of noise energy in high frequency component, i.e., the tails of the signal y_e, we use a second order polynomial model to fit the signal as

$$y_e = a_2 x^2 + b_2 x + c_2. \tag{10}$$

Then, by examining the convexity of the estimated signal, i.e., the sign of the estimated $\hat{a}_2$, we can detect whether the noise energy increases with $|x|$.

Combining these two features together, the decision rule of our detector is as follows,

$$\hat{H} = \begin{cases} H_0, & \text{if PSNR} < \tau_1 \text{ and } \hat{a}_2 < 0, \\ H_1, & \text{if PSNR} \ge \tau_2, \\ H_2, & \text{if PSNR} < \tau_1 \text{ and } \hat{a}_2 > 0, \\ H_3, & \text{if } \tau_1 \le \text{PSNR} < \tau_2 \text{ and } \hat{a}_2 < 0, \\ H_4, & \text{if } \tau_1 \le \text{PSNR} < \tau_2 \text{ and } \hat{a}_2 > 0. \end{cases} \tag{11}$$

The parameters of our proposed detector are $\underline{\theta} = (\tau_1, \tau_2)$.

5. SIMULATION RESULTS

Although detectors have been proposed to detect the order of resizing and blurring, there are cases where the fingerprints are so weak that we can hardly detect the order of these two operations, as we have shown in Fig. 2. Using our information theoretical framework and criteria, we can know when this order can or cannot be detected.

We first set up a testing image databased based on the 1338 images in the UCID database [20]. To simulate blurring operations, we used Gaussian blur with filter window 5 by 5 and variance ν. We used a wide range of resizing factors, denoted as s, from 0.5 to 2. For each $s = \{0.5, 0.55, ..., 2\}$ and $\nu = \{0.5, 0.55, ..., 1\}$, the testing database contains: 1338 unaltered images, 1338 resized images with scaling factor s, 1338 blurred images with Gaussian variance ν, 1338 blurred then resized images, and 1338 resized then blurred images. To avoid distorting the image too much, we set $\nu \le 1$, which was obtained by calculating the distortion introduced by blurring using the structure similarity (SSIM) index [21] and then setting the reasonable SSIM values to be greater than 0.9.

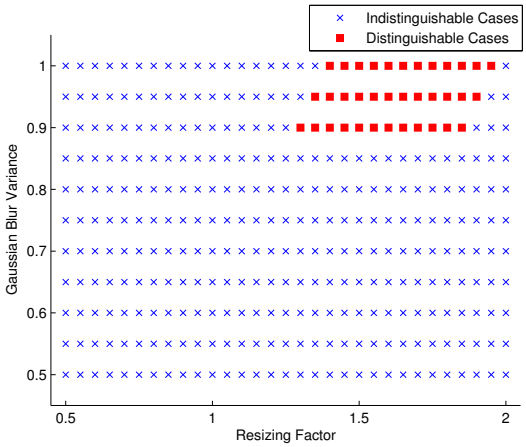

Fig. 4: Distinguishability test results of detecting the order of resizing and blurring by applying our information theoretical framework and criteria.

Based on the detector with tunable parameters in last section, we use our information theoretical framework and criteria to obtain the distinguishable and indistinguishable cases for different pairs of resizing factor and blurring parameter. Specifically, for each value of s and ν, we first calculate the transition probability matrices (2) for all possible values of detector parameters $\underline{\theta}^*$. Then, using definition 1, we can obtain the best parameters $\underline{\theta}^*$. Based on this best detector, we use definition 2 to determine whether we can or cannot distinguish all hypotheses in (1), i.e., can or cannot detect the order of resizing and blurring. Fig. 4 shows the results for all combinations of resizing factors and blurring parameters.

To understand why we cannot detect the order of resizing and blurring in those indistinguishable cases, we analyzed the indistinguishable cases close to the range of the distinguishable cases to find what makes the distinguishable cases become indistinguishable, i.e., which condition was violated in (4). We have found that, for most cases, the reason was that the hypothesis of resizing then blurring and blurring then resizing cannot be distinguished. This matches the example we have shown in Fig. 2 where the fingerprints in Fig. 2(b) and Fig. 2(c) are similar.

6. CONCLUSION

In this paper, we studied the fundamental question of when we can or cannot detect the order of operations. To answer this question, we formulated the order detection problem into a general multiple hypotheses testing problem. Then, we proposed an information theoretical framework and mutual information based criteria to obtain the best detector and answer the question of when we can distinguish all considered hypotheses. For case study, we examined the problem of detecting the order of resizing and blurring, where the order may not always be detectable. Forensic techniques have been proposed to detect the order of these two operations, and simulation results have shown the effectiveness of our proposed framework and criteria.

7. REFERENCES

[1] Z. Fan and R. L. de Queiroz, "Identification of bitmap compression history: JPEG detection and quantizer estimation," *IEEE Transactions on Image Processing*, vol. 12, no. 2, pp. 230 – 235, 2003.

[2] A. C. Popescu and H. Farid, "Statistical tools for digital forensics," in *6th International Workshop on Information Hiding*, Toronto, Canada, 2004.

[3] B. Li, Y. Q. Shi, and J. Huang, "Detecting doubly compressed jpeg images by using mode based first digit features," in *Multimedia Signal Processing, 2008 IEEE 10th Workshop on*, Oct 2008, pp. 730–735.

[4] A. C. Popescu and H. Farid, "Exposing digital forgeries by detecting traces of re-sampling," *IEEE Transactions on Signal Processing*, vol. 53, no. 2, pp. 758–767, Feb. 2005.

[5] M. Kirchner, "Fast and reliable resampling detection by spectral analysis of fixed linear predictor residue," in *Proceedings of the 10th ACM Workshop on Multimedia and Security*, New York, NY, USA, 2008, MM&Sec '08, pp. 11–20, ACM.

[6] M. C. Stamm and K. J. R. Liu, "Forensic detection of image manipulation using statistical intrinsic fingerprints," *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 3, pp. 492–506, Sep. 2010.

[7] H. Tong, M. Li, H. Zhang, and C. Zhang, "Blur detection for digital images using wavelet transform," in *Multimedia and Expo, 2004. ICME '04. 2004 IEEE International Conference on*, Jun. 2004, vol. 1, pp. 17–20 Vol.1.

[8] G. Cao, Y. Zhao, and R. Ni, "Edge-based blur metric for tamper detection," in *Journal of Information Hiding and Multimedia Signal Processing*, Jan. 2010, pp. 20–27.

[9] B. Su, S. Lu, and C. L. Tan, "Blurred image region detection and classification," in *Proceedings of the 19th ACM International Conference on Multimedia*, New York, 2011, MM '11, pp. 1397–1400, ACM.

[10] X. Chu, M. C. Stamm, and K. J. R. Liu, "Compressive sensing forensics," *IEEE Transactions on Information Forensics and Security*, , no. 99, Mar. 2015.

[11] M. C. Stamm, W. S. Lin, and K. J. R. Liu, "Temporal forensics and anti-forensics for motion compensated video," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 4, pp. 1315–1329, 2012.

[12] M. C. Stamm, M. Wu, and K. J. R. Liu, "Information forensics: An overview of the first decade," *IEEE Access*, vol. 1, pp. 167–200, 2013.

[13] T. Bianchi and A. Piva, "Reverse engineering of double jpeg compression in the presence of image resizing," in *IEEE International Workshop on Information Forensics and Security*, Dec. 2012, pp. 127–132.

[14] P. Ferrara, T. Bianchi, A. De Rosa, and A. Piva, "Reverse engineering of double compressed images in the presence of contrast enhancement," in *IEEE 15th International Workshop on Multimedia Signal Processing*, Sep. 2013, pp. 141–146.

[15] G. Cao, Y. Zhao, R. Ni, and X. Li, "Contrast enhancement-based forensics in digital images," *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 3, pp. 515–525, Mar. 2014.

[16] V. Conotter, P. Comesana, and F. Perez-Gonzalez, "Forensic detection of processing operator chains: recovering the history of filtered jpeg images," *Information Forensics and Security, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2015.

[17] M. C. Stamm, X. Chu, and K. J. R. Liu, "Forensically determining the order of signal processing operations," in *IEEE International Workshop on Information Forensics and Security*, Nov. 2013, pp. 162–167.

[18] P. Comesaña, "Detection and information theoretic measures for quantifying the distinguishability between multimedia operator chains," in *IEEE Workshop on Information Forensics and Security*, Tenerife, Spain, 2012.

[19] X. Chu, Y. Chen, M. C. Stamm, and K. J. R. Liu, "Information theoretical limit of compression forensics," in *IEEE International Conference on Acoustics, Speech and Signal Processing*, May 2014, pp. 2689–2693.

[20] G. Schaefer and M. Stich, "UCID: An uncompressed color image database," *Proc. SPIE: Storage and Retrieval Methods and Applications for Multimedia*, vol. 5307, pp. 472480, 2004.

[21] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004.