Journal of Electronic Imaging

SPIEDigitalLibrary.org/jei

Content authentication for printed images utilizing high capacity data hiding

Orhan Bulan
Gaurav Sharma

SPIE

IS&T
imaging.org

Content authentication for printed images utilizing high capacity data hiding

Orhan Bulan
University of Rochester
Department of Electrical and Computer Engineering
Rochester, New York 14627-0126
E-mail: bulanorhan@gmail.com

Gaurav Sharma
University of Rochester
Department of Electrical and Computer Engineering
Rochester, New York 14627-0126
and
University of Rochester
Medical Center
Department of Biostatistics and Computational Biology
Rochester, New York 14630

Abstract. *Authentication of content in printed images poses a challenge that cannot be addressed by conventional digital signature schemes because under the analog transport provided by the printing channel the verifier does not have access to the original digital content in pristine form. We present a method for cryptography-based authentication of the content in printed images that also provides the capability for identifying localized changes made by informed malicious attackers—key functionality that is missing in print scan robust hashes that have traditionally been used for print content authentication. The proposed method operates by embedding, within the printed image, an authentication signature that consists of an encrypted thumbnail of the image using a high capacity data hiding method for halftone images. To authenticate the content, the embedded signature is extracted from a scan of the printed image and, after decryption, compared with the printed content. An implementation of the method that incorporates human or automated verification and identifies potential local tampering by informed malicious attackers is developed and successfully demonstrated.* © *2013 SPIE and IS&T* [DOI: 10.1117/1.JEI.22.3.033006]

1 Introduction

Printed images are extensively used in official identification documents (e.g., passports, driving licenses, and other identification documents), as prosecutorial legal evidence in a court of law, and for communicating sensitive military information. In these applications, verification of the image authenticity is highly desirable to confirm the information conveyed by the image before taking an action whose consequences may be serious and costly.

Traditional approaches for establishing the integrity of a printed document rely on specialty papers,[1,2] customized inks,[3,4] special optical,[5] or print reproduction techniques.[6–10] These techniques provide protection against counterfeiting and their security is derived from the fact that the special materials and high precision reproduction equipment and techniques are not accessible to most users. More recently, alternative techniques enabled by digital processing capabilities have also been proposed for anti-counterfeiting applications, where random, hard to reproduce, characteristics of a specific print are used as a security feature.[11,–14] Some of these digital techniques[12,13] have also been augmented to provide authentication of printed text content, usually by including an encrypted version of the text as a barcode within the print. Authentication of image content within printed documents has, however, received only limited attention.[15] A recent incident[16] in which the image content was manipulated while retaining the print substrate, further highlights the need for such authentication.

For digital images, digital signatures enable content authentication. Digital signatures rely on asymmetric key cryptography where a user has a secret private key, known only to him/her-self, and a public key that is uniquely associated with the user and available to anyone through a suitable public key infrastructure (PKI).[17] To sign a block of data, the user uses his/her private key to compute a key-dependent cryptographic hash of the data, which takes the form of a short binary data string. The signature is distributed along with the data. Using the signing user's public key with the data anyone can validate the signature. Since, the signature has the property that it is computationally infeasible (Strictly speaking, for cryptographic constructions of existing digital signatures, the computational infeasibility is believed to hold, though this has not been proven) for an adversary to compute the signing user's signature for a given data block without access to the private key, a digital signature cannot be faked by unauthorized individuals. Furthermore, unless the private key is compromised, a valid signature uniquely associates the signed data with the signing entity, providing

Paper 12477 received Nov. 18, 2012; revised manuscript received May 16, 2013; accepted for publication Jun. 24, 2013; published online Jul. 26, 2013.

nonrepudiability. For image authentication, signature schemes are usually designed to provide tamper localization by computing multiple signatures corresponding to localized regions of the image. These signatures can be conveniently embedded in the image itself by using data hiding techniques, resulting in what are commonly referred to as fragile authentication watermarks.[18,19]

A key feature of digital signatures is that they are extremely sensitive to changes in the data being signed; a single bit modification in the data typically changes the signature dramatically.[17] While this contributes to the security of the digital signature schemes, it poses a limitation when applied to image authentication because images are regarded as authentic if the perceived content is unchanged, even though the exact digital data representing the image may be altered. For applications where authentication is desired in the presence of benign image modifications, robust hashes have, therefore, been proposed.[20–23] For the purpose of our description, these robust hashes can be assumed to consist of a private-key encrypted version of the key image features, such as corner points and edges that are robust to incidental modifications such as JPEG compression, blurring, or scaling, but sensitive to malicious modification of the content. The authenticity of a test image is established by computing the features from the image and comparing if these are within a small tolerance of the decrypted image features from the purported robust hash.

For printed images, the signature based digital image authentication framework is inapplicable because the printing process constitutes an analog channel that makes digital signatures ineffective; due to perturbations introduced in the printing process, the original digital image cannot be recovered from the print and therefore, a corresponding signature cannot be obtained for authentication. Robust hashes for digital image authentication also, typically, have limited effectiveness for printed images because defining feature points that are fragile to malicious manipulation of the content but robust to the severe distortions introduced by printing process, is a difficult problem. In particular, printing processes commonly employ a binary halftone representation which matches the original image only in the low frequency representation perceived by the human visual system.[24] Common feature detectors yield a large number of spurious features when applied directly to scans of printed halftone images. De-screened versions of images, obtained by low-pass filtering to eliminate the halftone structure, also cause unacceptable loss of features. Finally, the presence of locally varying geometric distortion[25,26] introduced in the printing process also poses a significant challenge for robust hashes.

To address these challenges, specialized robust hashes have been proposed explicitly for printed image content authentication.[27–29] In this paper, we propose and demonstrate a method for content authentication of printed images with enhanced functionality compared with robust hashes. Prior to printing, we compute a compressed, scaled, and encrypted a thumbnail of the digital image as a signature of the print content. This signature is then embedded within the printed image using a high capacity data hiding technique that is resilient to the distortions in the print-scan process—our specific implementation uses our recently proposed data hiding method in halftone images.[30] At the receiver, first the embedded data is extracted and decrypted to obtain the embedded thumbnail. The authenticity of the image is then verified by comparing the print content against the reference thumbnail extracted from the printed image. Visual comparison is augmented by adding an automated procedure for verification and identification of local tampering by informed malicious attackers. The implementation is validated via experiments that characterize the empirical performance of the scheme. In addition to cryptography-based authentication of the printed content our scheme enables identification and localization of changes and provides indication of the original content in the presence of such manipulations, something that has not hitherto been feasible. The comprehensive approach to content authentication enabled by our proposed scheme also closes potential security loopholes that exist in several of the print-scan robust hashes that ignore the consideration of local manipulation by informed malicious attackers.

This paper is organized as follows. In Sec. 2, we present the proposed method for content authentication of printed images and consider two possible attacks against the method. A specific implementation of the proposed method is described in Sec. 3. Section 4 presents experimental validation of the method and a limited comparison against robust hashing based alternatives. Section 5 concludes the paper with a summary of the main findings, a discussion of relative merits and limitations of the proposed scheme and alternative constructions based on existing print scan robust hashes, and a commentary on open problems in printed document security. Appendix A summarizes the data embedding and extraction scheme that we use in our implementation. A preliminary version of part of the work in this paper has been previously presented in Ref. 31.

2 Content Authentication for Printed Images

Figure 1 is the overall block diagram of the proposed method that shows both signature embedding and verification phases. The signature embedding phase has two steps: signature generation and signature transport, where the latter refers to the embedding of the signature data within the printed image. As a signature, we use an encrypted digital thumbnail of the original image, obtained by using either symmetric or asymmetric key cryptography. The computed signature data is embedded within the printed image by a data hiding technique that is robust to distortions inherent in the print-scan process. Because the thumbnail is larger than typical digital signatures, the data hiding technique has to offer high capacity in order to carry the signature within the print content.

The verification phase of the proposed method is illustrated in Fig. 1(b). Data is first extracted from a scan of a printed image under consideration. The extracted data is decrypted by using the key for the entity or individual that is claimed to be the signatory, to obtain (what is presumed to be) a thumbnail. Authenticity is then established by comparing the extracted thumbnail against the printed/scanned image content. If the printed image content is indeed signed by the claimed signatory and the embedded digital data is successfully recovered from the printed image, the extracted thumbnail will match the contents of the printed image indicating authenticity. The use of an encrypted digital thumbnail as a signature provides considerable flexibility and ease of use; the matching step can be either visual or automatic.

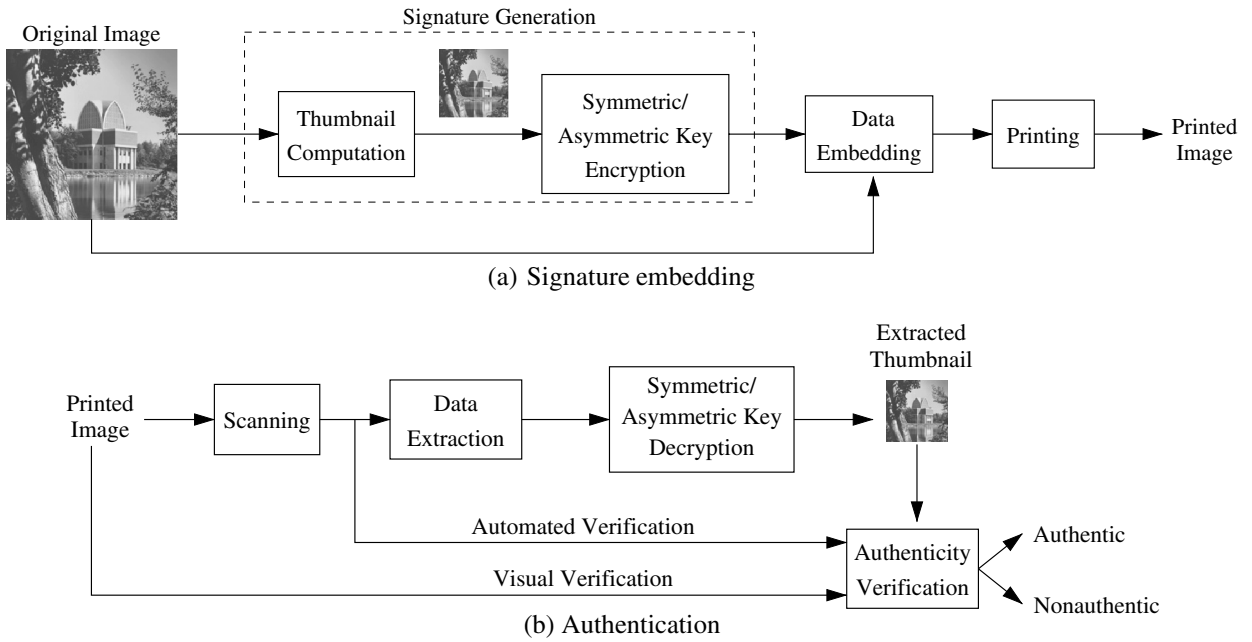

Fig. 1 Content authentication framework for printed images: (a) signature embedding and (b) authentication verification.

For visual verification, a human observer is interjected to make a decision about authenticity by comparing the extracted thumbnail against the printed content. Automated verification, on the other hand, compares the extracted thumbnail against the scanned image without involving a human observer, attempting to once again make a decision regarding authenticity and tamper localization.

For the computation and verification of signatures, one can use either symmetric[17] or asymmetric[17] key cryptography. The latter has the advantage of allowing public authentication of the printed images by exploiting PKI in much the same way as it is utilized for digital signatures.[17] The security of authentication is cryptographically ensured and does not rely on secrecy of the technique used for signature computation or data embedding. The resulting method, therefore, also inherits several of the advantages inherent in modern cryptography, including nonrepudiability of signatures and the ability to revoke compromised keys and authorize new ones as required.

When details of the data embedding method are public, the potential for malicious attacks needs to be considered. The use of cryptography provides security against unauthorized users. Users without access to the cryptographic key used in embedding cannot (Strictly speaking, "cannot" here means that unless the cryptographic scheme is compromised the probability of success per attempt is 2^{-N_K} where N_K is the number of bits in the key) generate the signature for a nonauthentic image. An unauthorized user can embed data in a printed image, possibly even an encrypted thumbnail obtained with a random key, and the embedded data will be extracted correctly during the validation process. Upon decryption, however, instead of the thumbnail of the printed image, a random bitstream is obtained as shown in Fig. 2(a), which indicates that the printed image content is nonauthentic. The fact that the signature consists of a thumbnail of the image content provides security against transfer attacks where a malicious user can extract the digital (encrypted) data from an authentic printed image and embed this data within another image, which is a specific instance of the more broadly applicable watermark copy attack.[32] In this scenario, as shown in Fig. 2(b), a thumbnail is successfully extracted during the authentication step but it does not match the printed image content, again indicating nonauthenticity. Malicious attacks may also be designed more intelligently than the two scenarios considered earlier in this paragraph. A particularly pernicious attack, which we refer to as the local manipulation and transfer attack, combines scanning and reprinting of an image accompanied by a transfer of the original (encrypted) signature data to the new print. In this case, a valid thumbnail is recovered and matches significant parts of the overall content of the printed image. A comparison of the extracted thumbnail against the printed image can, however, still establish the nonauthenticity and also localize tampered regions.

3 Implementation

The method for content authentication of printed images proposed in Sec. 2 can be implemented using different techniques for thumbnail computation, encryption, and data embedding, as long as the requirements for individual system elements are satisfied. In this section, we describe a specific implementation of the proposed framework. Where appropriate, we also outline a few other potential alternatives and their advantages.

3.1 *Signature Computation*

In our specific implementation, a signature is obtained by JPEG compressing a scaled down version of the image to obtain a thumbnail which is then encrypted by using the Advanced Encryption Standard (AES).[33] Though our specific construction here is based on a symmetric key protocol, asymmetric key cryptography can be utilized equally readily, with its attendant advantage of private signing with public authentication capability. The higher computational complexity of asymmetric key cryptography may be mitigated

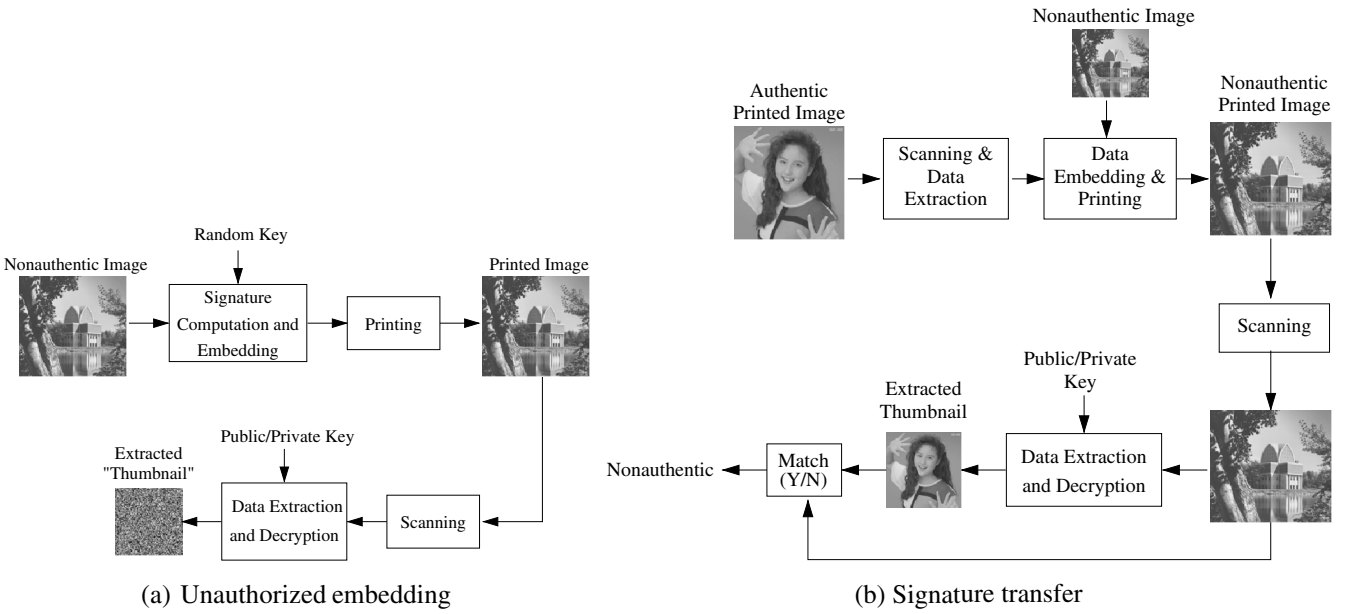

Fig. 2 Detection of nonauthentic content under (a) unauthorized embedding and (b) signature transfer.

in this setting by using a digitally signed version of the thumbnail as a signature instead of a directly encrypted version. For thumbnail generation we can also use JPEG2000,[34,35] which offers a better rate-distortion performance.

3.2 *Signature Transport*

The proposed framework requires that the computed signature be embedded in the printed image using a suitable data hiding technique. This requirement imposes two constraints on the data hiding technique: high capacity and robustness against the print-scan process. The latter requirement is obvious from the problem setting and high capacity is necessary because the encrypted thumbnail is a rather large size signature that must be embedded in the print content with minimal embedding distortion. Any hardcopy data hiding technique satisfying these conditions can be used in the proposed method, such as Refs. 30 and 36–39. Note that these high capacity data hiding techniques are generally print technology dependent. Our specific implementation uses halftone-dot orientation modulation[30] to embed data in clustered-dot halftones commonly used in laser printers. For ink jet printing, alternative high capacity data hiding techniques[36,38] that are well suited for error diffusion halftones would be more appropriate. A brief summary of our data embedding and extraction technique is included in Appendix A to facilitate understanding of this paper by itself and to specifically describe how the printed image is registered to the thumbnail in the signature, a step that is crucial for the automated verification step in Sec. 3.3.

3.3 *Signature-Based Verification*

The data extracted from the print is decrypted by using the AES algorithm with the key of the presumed signatory at the embedder. If the keys used at the embedder and receiver are the same and the image is not manipulated in the channel, this process recovers the digital thumbnail of the printed image. If a decoding failure occurs due to uncorrectable errors in the data extraction process, the print is deemed nonauthentic. This is desirable in authentication applications because a far greater cost is usually associated with false positives (i.e., declaring a nonauthentic image as authentic) than with false negatives (i.e., declaring an authentic image as nonauthentic), where one may, for instance, resort to alternate authentication mechanisms. Similarly, if the print in question does not carry any data in itself, error correction decoding fails in the data extraction phase and the image is declared nonauthentic.

Once data is successfully recovered by the error correction decoding process, authenticity of the print image can then be established by visual/automated comparison against the extracted thumbnail. Figure 3 illustrates a specific implementation of the automated verification procedure for the proposed method. Using the synchronization information obtained during the data embedding and extraction phase, global and local geometric distortion introduced in the print-scan process is compensated for, and the scanned image is scaled down to the extracted thumbnail size, assumed to be $M \times N$ pixels in the following description. The authenticity of the content is then established in two steps. In the first step, the extracted thumbnail is correlated with the processed image to identify situations where the thumbnail is entirely unrelated to the content of the printed image. For this purpose, the correlation between the extracted thumbnail A and the preprocessed scanned image B is calculated as

$$\rho(A,B) = \frac{\sum_{i=1}^{M}\sum_{j=1}^{N}[A(i,j)-\bar{A}][B(i,j)-\bar{B}]}{\sqrt{\{\sum_{i=1}^{M}\sum_{j=1}^{N}[A(i,j)-\bar{A}]^2\}\{\sum_{i=1}^{M}\sum_{j=1}^{N}[B(i,j)-\bar{B}]^2\}}}, \quad (1)$$

where $\bar{A}$ and $\bar{B}$ are the mean of the thumbnail and the scanned image and i and j are the pixel position indices. If the correlation is smaller than a predefined threshold value τ, the image is deemed nonauthentic. High correlation,

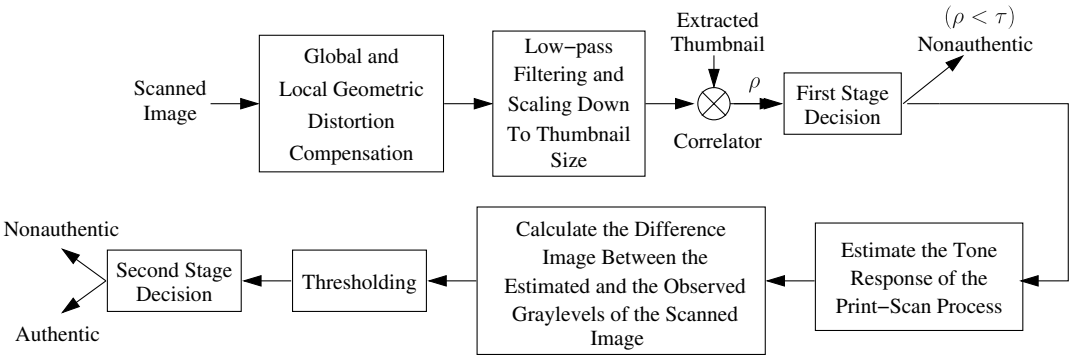

Fig. 3 Automated verification.

however, does not always guarantee authenticity. Specifically, under the local manipulation and transfer attack outlined in Sec. 2, the global similarity between the images can result in a high correlation with the extracted thumbnail despite local manipulations. In order to identify nonauthentic regions in these cases, we first model graylevel transformation of the print-scan process by two parameters as graylevel scaling and offset based on the observed experimental data. We next estimate these parameters by using the least squares (LS) approach. The LS problem is formulated as

$$\min_{a,b}(a\mathbf{x}+b\mathbf{1}-\mathbf{y})^T(a\mathbf{x}+b\mathbf{1}-\mathbf{y}), \tag{2}$$

where $\mathbf{x}$ and $\mathbf{y}$ are $N\times 1$ vectors containing the graylevels of the extracted thumbnail and the corresponding graylevels observed in the scanned image, respectively, and $\mathbf{1}$ represents $N\times 1$ vector whose entries are all ones. We estimate the graylevel scaling and offset parameters a and b by solving the unconstrained optimization problem as

$$[\,a \quad b\,]^T=(\mathbf{C}^T\mathbf{C})^{-1}\mathbf{C}^T\mathbf{y}, \tag{3}$$

where $\mathbf{C}=[\,\mathbf{x} \quad \mathbf{1}\,]$. After estimating the graylevel scaling and offset parameters of the print-scan process, we calculate the difference between the estimated and the observed graylevels of the scanned image. The estimated graylevels closely follow the observed graylevels of the scanned image in regions where the scanned image and the extracted thumbnail match. In the manipulated regions, however, observed graylevels deviate from the estimated values since the graylevels observed in these regions do not fit in the graylevel transformation model of the print-scan process. The manipulated regions in the scanned image are then detected by thresholding the difference image between the estimated and the observed graylevels of the scanned image to establish authenticity and identify regions of tampering.

4 Experimental Results

We evaluated the performance of our implementation using an experimental setup consisting of an electrophotographic printer with an addressability of 2400 dots per inch (dpi) and a desktop scanner with a 1200 dpi resolution. We tested the method using 30 grayscale images with varying content; thumbnail versions of 16 of these images are shown in Fig. 4. The halftone-dot orientation modulation method that we used in our implementation for data hiding in printed images provides, in conjunction with the error correction encoding, the capability to communicate 12.6 Kbytes within each image when printed on an 8×8 in^2 footprint. We obtain thumbnails of the images that fit within this estimated error free operational rate by scaling down and compressing the images with JPEG quality factor of 80. The size of the thumbnails varies depending on the image content. As described in Sec. 3, the AES encrypted thumbnail is encoded with a rate $1/4$ repeat-accumulate (RA) encoder and embedded within the printed image during the halftoning process by using the orientation modulation embedding method. For binary modulation, an RA code rate $1/4$ is adequate across images with a wide variation in image content, as is validated by the results we present here and additional results in Ref. 30, where the error-free communication rate was examined (without consideration of the authentication application considered here).

Printed images are authenticated using the method outlined in Sec. 3. We first consider the scenario where the print under consideration corresponds to a signed image that has not been manipulated or modified. The images in Fig. 4(a)–4(p) show the thumbnails obtained in the validation phase for 16 of the images. In each case, the signature data was successfully recovered by the validation process. It can be seen that the thumbnails provide an adequate representation for visual verification of the authenticity of the printed content and also allow for the viewer to identify potential local manipulations induced either by modifying the print or by the local manipulation and transfer attack described in Sec. 2.

For benchmarking purposes, we also tested the robust hash based methods of Refs. 27 and 28 using the original and scanned images in our experimental setting, where in both cases we assumed perfect transport of the robust hashes via the data hiding scheme and evaluated the schemes with respect to their ability to withstand the print-scan distortion in our experimental setup and performance under local modification. For the method in Ref. 28, we used default parameters for feature extraction and hash matching and scaled the scanned images to match the original image sizes prior to robust hash computation. (The method is sensitive to scaling and identifies all the scanned images as nonauthentic when the scaling does not match the original image sizes.) In this case, nine out of 30 images (i.e., 30% of all images) are declared as nonauthentic, where, in fact, all the images are authentic. This high false alarm rate is observed because distortions in the print-scan process (e.g., halftoning, local geometric distortions, noise in the graylevel) introduce

Fig. 4 Thumbnails corresponding to the digital data recovered from printed images. All thumbnails are compressed with a JPEG quality factor of 80. The size of the thumbnails varies depending on image content.

several spurious feature points and eliminate some of the feature points detected in the original image. The robust hash based method presented in Ref. 27 was also evaluated using recommended parameter settings. The original and scanned images were scaled (The scaling eliminates the halftone structure in the scanned image and also accounts for loss of resolution in the printing process. Better performance is obtained with the scaling compared with the alternative of working with the native scan resolution), to a size of 256×384 and divided into a 2×3 array of nonoverlapping 128×128 pixel blocks. A 128×128 two-dimensional (2-D) discrete cosine transform (DCT) was computed for each block from which a block signature consisting of an 16×16 binary string was computed by adaptively thresholding selected DCT coefficients as described in Ref. 27. The signature for the whole image was obtained by concatenating together the signatures for the individual blocks and Hamming distances between the signatures for the original and, as in Ref. 27, the images obtained after the print-scan process were used for authenticating the image content. Figure 5(a) shows (on a single plot) two normalized histograms of Hamming distances between the signatures calculated from a scan of a printed image and a test digital image, where the histogram labeled "authentic" corresponds to the case when the test digital image corresponds to the original from which the print was generated and the histogram labeled "nonauthentic" represents the case when the test digital image content is different from the printed image. The two histograms are well separated in Fig. 5(a) and, therefore, in the absence of local manipulation, there is ample latitude to reliably distinguish authentic images from nonauthentic images based on the signature scheme of Ref. 27. The performance of this scheme, however, deteriorates significantly under local manipulation, a fact that can be seen from Fig. 5(b), where we have repeated the histograms shown in Fig. 5(a), but superimposed on the same figure the normalized histogram of

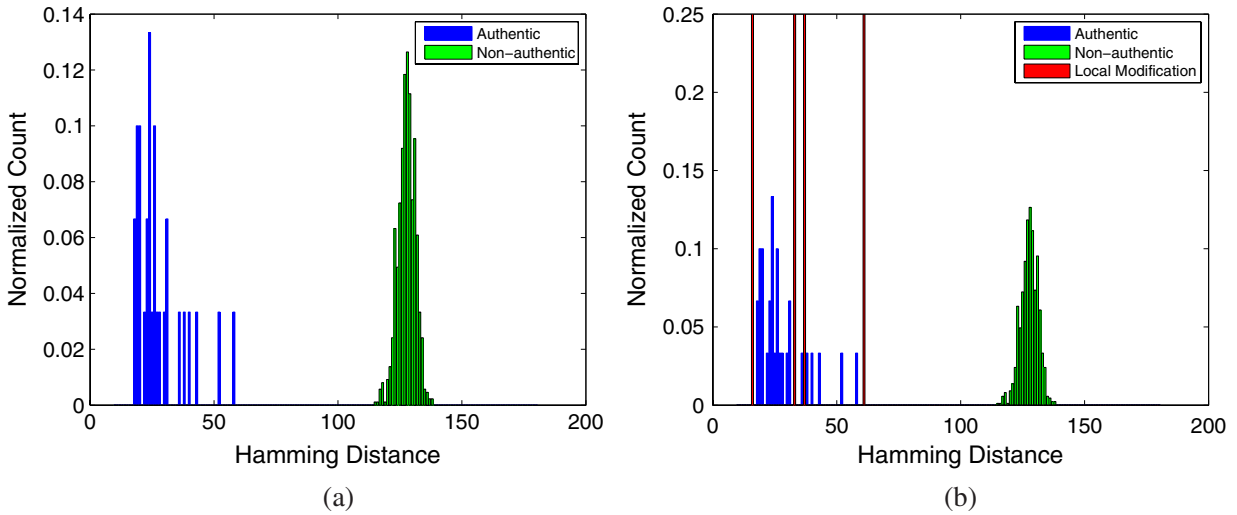

Fig. 5 Normalized histograms of Hamming distances between the signatures calculated from a scanned image and a digital test image using the method proposed in Ref. 27. Histograms labeled "authentic," "nonauthentic" correspond, respectively, to situations where the digital test image corresponds to the original image from which the print was generated and an image with different content. The histogram labeled "local modification" corresponds to the situation where the scanned image corresponds to a print generated from the digital test image but subjected to local modification.

Hamming distances between the signatures of the locally manipulated images in our experiments (see Sec. 4.2) and the corresponding premanipulation original digital images used for printing. The histogram bins with nonzero values for our locally manipulated images lie in the same range as the peaks for the "authentic" scenario and thus, these images would be incorrectly classified as "authentic". These results also lead to the following important conclusion: robust hash based print scan authentication schemes that do not consider the local manipulation and transfer attack potentially have a security loophole and images they deem "authentic" may potentially be locally manipulated.

We next consider automated verification of authenticity, focusing particularly on the performance of the method in detecting nonauthentic images and tamper localization. For this purpose, we consider three print scenarios with nonauthentic content described in Sec. 2 as unauthorized users, signature transfer, and local manipulation followed by signature transfer.

4.1 *Unauthorized Users/Signature Transfer*

When the printed image under consideration is generated by an unauthorized user that does not have access to the private key of the signing entity, either a decoding failure occurs in the error correction decoder (i.e., corresponding to the case when no data is embedded in the printed image) or a random bit stream is extracted resembling random noise (i.e., when an encrypted thumbnail is embedded within the printed image obtained by a random key). In the event of a decoding failure, the print content is declared as nonauthentic and no further analysis is conducted. When decoding succeeds and a random bit stream is recovered, invariably it does not corresponds to a valid thumbnail file format. To be conservative in our analysis, we, however, assume that in this case a pseudorandom thumbnail is recovered. The first stage of the automated verification process in Sec. 3.3 corresponding to the correlation detector is then invoked. Depending on the threshold value (τ) selected, in this first stage, the authentication system has two types of errors: (a) false positives, i.e., situations where a nonauthentic image is deemed authentic and (b) false negatives, i.e., situations when an authentic image is declared as nonauthentic. The probability of these errors is, respectively,

$$P_{\mathrm{FP}} = Pr[\rho(A,B) > \tau | \text{image is nonauthentic}], \quad (4)$$

$$P_{\mathrm{FN}} = Pr[\rho(A,B) < \tau | \text{image is authentic}]. \quad (5)$$

We evaluate these two types of errors by observing the correlations obtained from the 30 images with varying content. For the authentic scenario, we calculate the correlations between preprocessed scanned images and their extracted thumbnails. We also calculate the correlation between the preprocessed scanned images and the thumbnail size images that resemble random noise for the unauthorized embedding case. The histogram of the resulting correlations is shown in Fig. 6. If the thumbnails are similar to random noise, the correlation values are accumulated around 0 whereas the correlations fluctuate around 0.95 when the thumbnails match with the print content. Based on the observed histograms in the figure, we model the correlations corresponding to authorized and unauthorized embedding with Gaussian distributions $\mathcal{N}(\mu,\sigma)$ and estimate the parameters of the distributions via a maximum likelihood estimator. P_{FP} and P_{FN} can be then estimated as

$$P_{\mathrm{FP}} = Q\left(\frac{\tau-\mu_2}{\sigma_2}\right), \qquad P_{\mathrm{FN}} = 1 - Q\left(\frac{\tau-\mu_1}{\sigma_1}\right), \quad (6)$$

where μ_1 and σ_1 are the mean and standard deviation of the distribution corresponding to nonauthentic case, μ_2 and σ_2 are the parameters for the authentic case, and $Q(\cdot)$ denotes the Q-function that evaluates the cumulative tail probability for a zero-mean, unit-variance Gaussian random variable. The variation of the probability of false positive and negative, i.e., P_{FP} and P_{FN}, across different threshold values is shown in Fig. 7. For a wide range of threshold values (0.1 to 0.85),

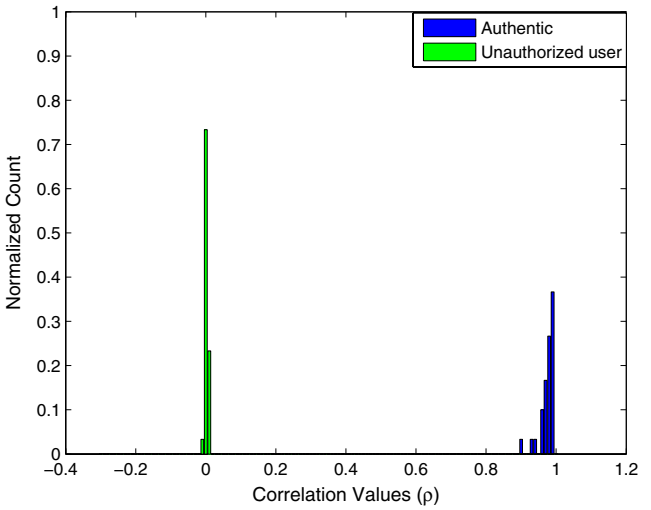

Fig. 6 Histogram of correlations of scanned images with extracted thumbnails under authentic and nonauthentic content. The cluster of histogram peaks in the vicinity of the value 1 represent correlations for images with authentic content where the thumbnail matches the image and the cluster of histogram peaks around the value 0 represents images with nonauthentic content where the extracted thumbnail corresponds to random noise.

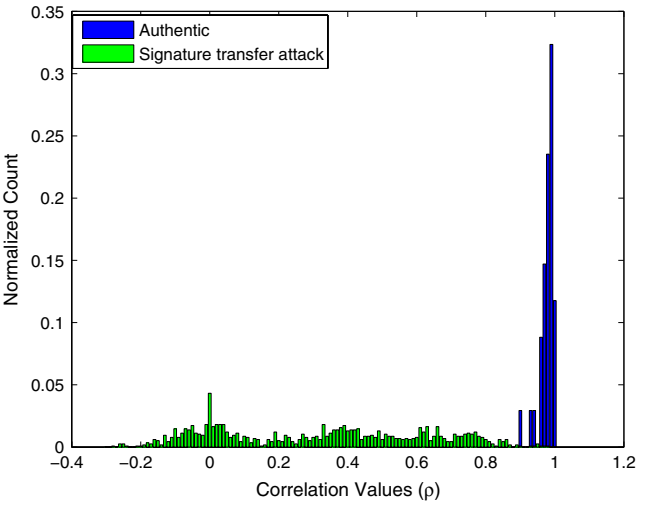

Fig. 8 Histogram of correlations of scanned images with extracted thumbnails under signature transfer attack. The cluster of histogram peaks in the vicinity of the value 1 represents correlations for images with authentic content where the thumbnail matches the image and the cluster of histogram peaks around the value 0 represents images with nonauthentic content where the extracted thumbnail corresponds to random noise.

both the probability of false positive (P_{FP}) and false negative (P_{FN}) is negligible ($<10^{-8}$), providing significant latitude in choice of the threshold value τ.

We next consider the performance of the proposed method in detection of a nonauthentic print generated by a signature transfer attack. For this purpose, we model the distribution of the correlations based on the observed histograms obtained from the experimental data. To obtain the statistics of the correlations, we repeat the experiment described in the preceding section (with the same set of images). After scanning the printed images, we calculate the correlations between the preprocessed scanned images and the thumbnails. Figure 8 shows the histogram of the resulting correlations. When the thumbnail matches the print content, the resulting correlation typically varies around 0.95. If the thumbnail is different from the print content, the correlation is almost uniformly distributed between -0.2 and 0.85, depend on the similarity between the thumbnail and the image content. Based on the observed histograms, we model the distribution of the correlations for the nonauthentic case with a uniform distribution and estimate P_{FP} and P_{FN} as

$$P_{\mathrm{FP}} = \frac{0.85 - \min(\tau, 0.85)}{1.05}, \quad P_{\mathrm{FN}} = 1 - Q\left(\frac{\tau - \mu_1}{\sigma_1}\right). \quad (7)$$

Note that P_{FN} remains the same as in Eq. (6) since the distribution of the correlations for the authentic case does not change. Figure 9 shows the projected receiver operating characteristic curve based on the estimated error probabilities P_{FP} and P_{FN}. Setting the threshold value $\tau = 0.85$ yields a false negative probability $P_{\mathrm{FN}} \approx 3 \times 10^{-9}$ whereas $P_{\mathrm{FP}} = 0$.

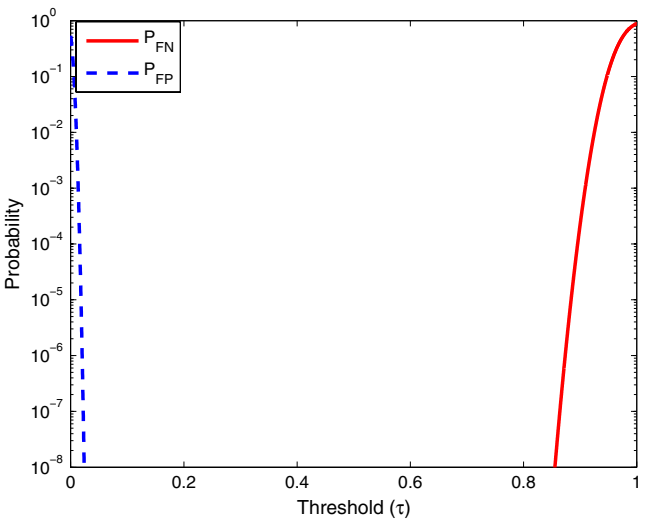

Fig. 7 Variation of modeled (first verification stage) false positive (P_{FP}) and false negative (P_{FN}) probabilities across different choices of the validation threshold τ.

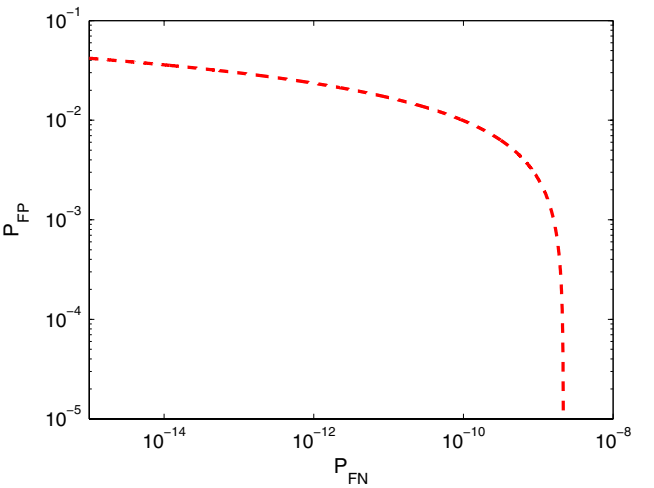

Fig. 9 The estimated receiver operating characteristic curve of the first stage of the automated verification in the case of a signature transfer attack.

4.2 *Local Manipulation and Signature Transfer*

We finally evaluate the performance of the proposed method for tamper detection when a nonauthentic printed image is generated by an intelligent malicious attacker using a local manipulation and signature transfer attack. The nonauthentic printed image in this scenario as well as the locally manipulated regions are detected in the second stage of the automated verification by estimating the parameters of the graylevel transformation of the print-scan process as described in Sec. 3.3. We illustrate the performance of the proposed method across the locally tampered images is shown in Fig. 10(a). Figure 10(b) shows the extracted thumbnails from the manipulated images and Fig. 10(c) shows the graylevels of the thumbnails and the preprocessed scanned images and the estimated graylevel transformation of the print-scan process for each image. As shown in the figure,

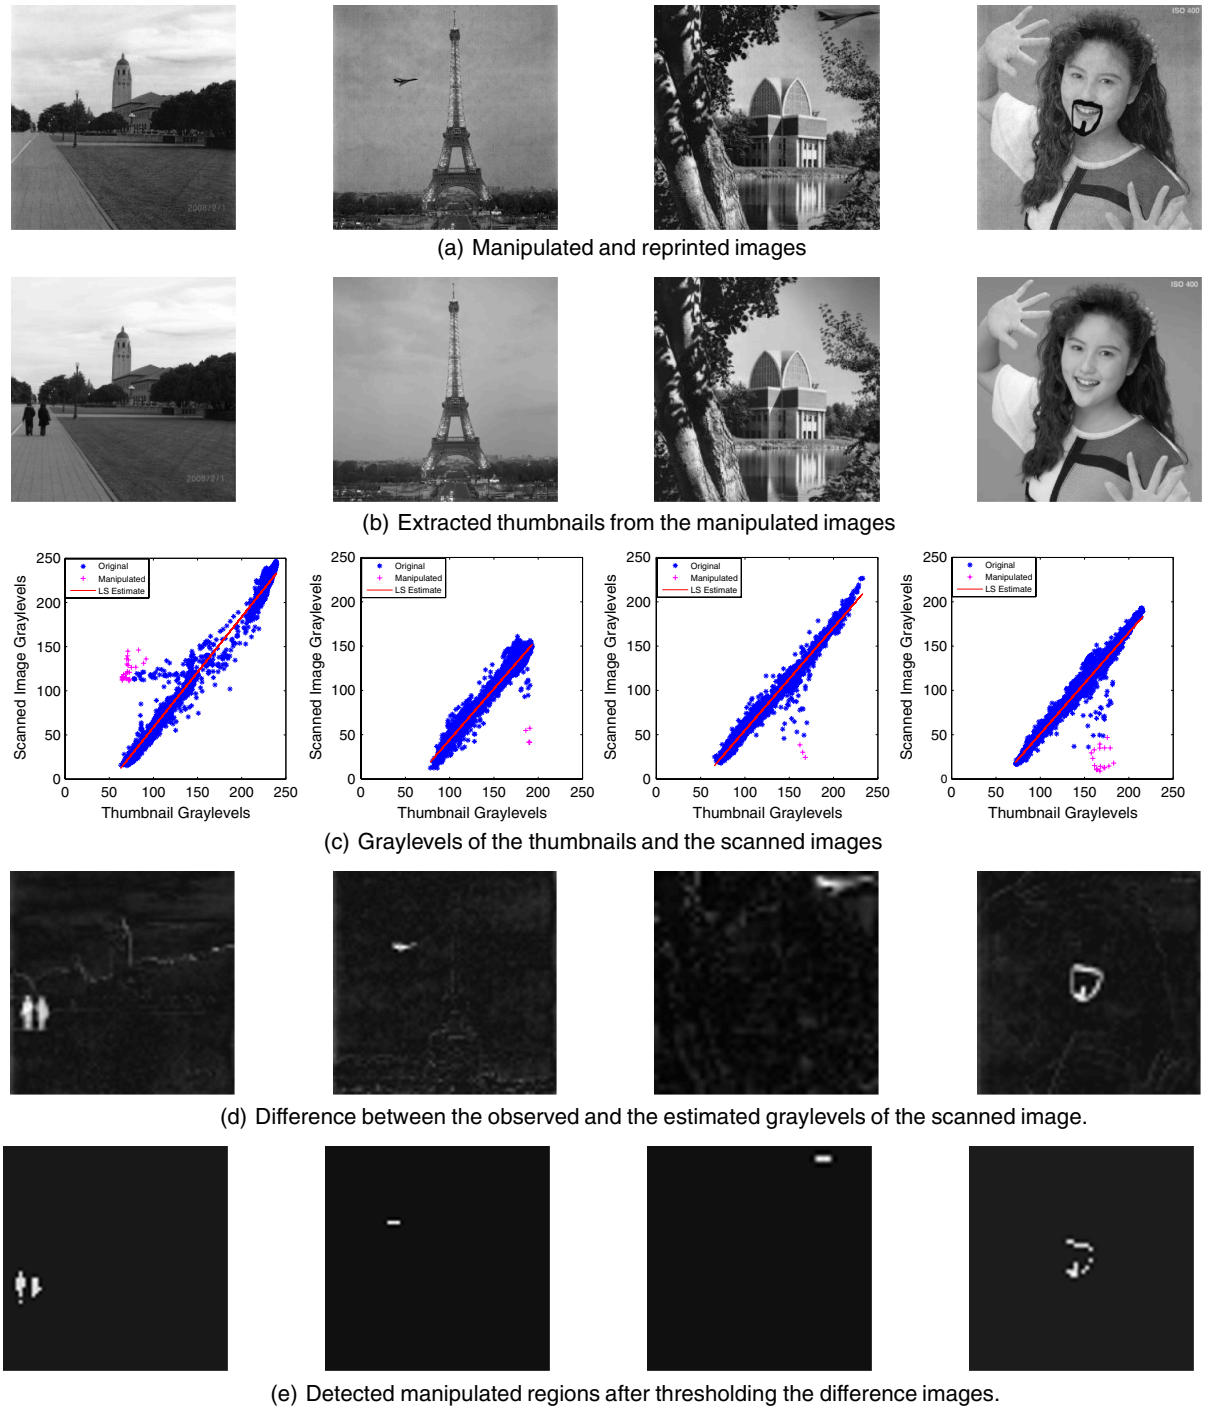

(a) Manipulated and reprinted images

(b) Extracted thumbnails from the manipulated images

(c) Graylevels of the thumbnails and the scanned images

(d) Difference between the observed and the estimated graylevels of the scanned image.

(e) Detected manipulated regions after thresholding the difference images.

Fig. 10 Detection of the manipulated regions when a nonauthentic printed image is generated by local manipulation and signature transfer attack. Manipulations are performed in digital domain after scanning the original printed image. Images in the second row are the thumbnails extracted from the manipulated images.

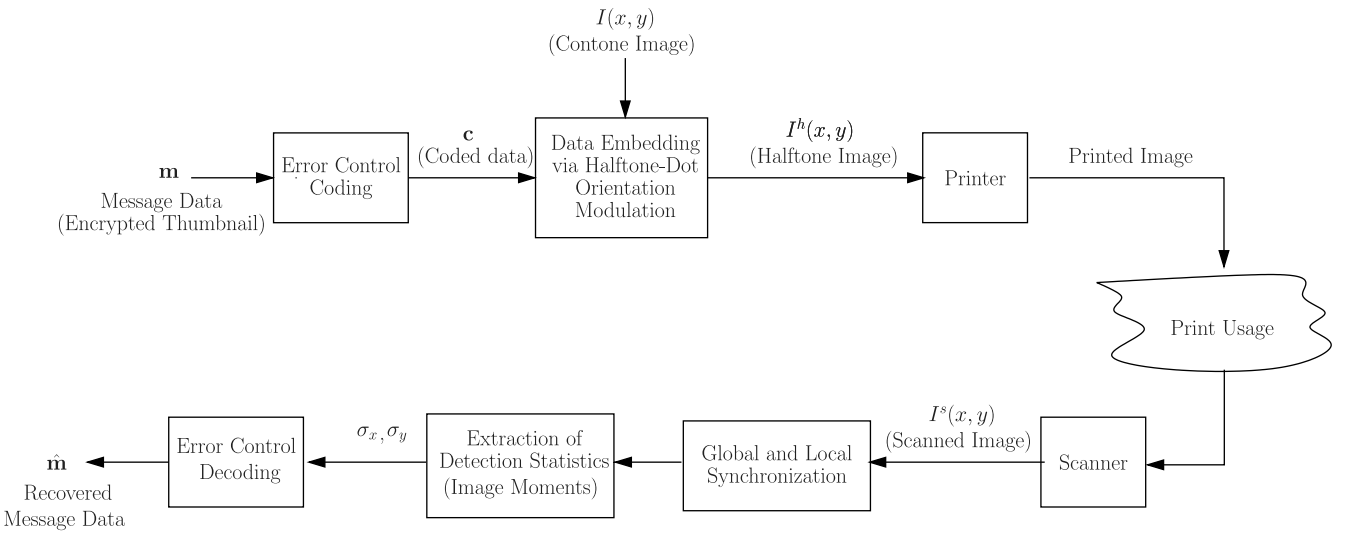

Fig. 11 Signature transport using orientation modulation based data hiding.

the graylevels of the manipulated regions do not fit the estimated model. For each 5×5 pixel region of the preprocessed scanned image, we then calculate the difference between the estimated and the observed graylevels of the preprocessed scanned image. The set of pixels in the difference image that deviate from the modeled transformation by more than $1.6\sigma_{\text{auth}}$ is shown in Fig. 10(e), where σ_{auth} is the root-mean squared deviation from the model observed for authentic images. The detected pixels can be seen to lie within the manipulated regions of Fig. 10(b). Note that an attack that transfers the signature from one image to another image with similar pictorial content, can also produce reasonably high correlation in the first stage of the detection process but is detected in the second stage.

5 Conclusion and Discussion

We develop a cryptography-based method for authenticating the content of printed images and demonstrate a successful experimental implementation of the proposed method. Our scheme improves upon the print content authentication functionality provided by print-scan robust hashes by enabling change localization and indication of the original content under informed malicious watermark transfer and local manipulation attacks.

To enable the functionality enhancements over robust hashes, our proposed print content authentication scheme requires a significantly higher capacity data hiding method. As an alternative, one could also consider enhancing the functionality of the robust hash based print content authentication schemes to allow localization of changes by using localized

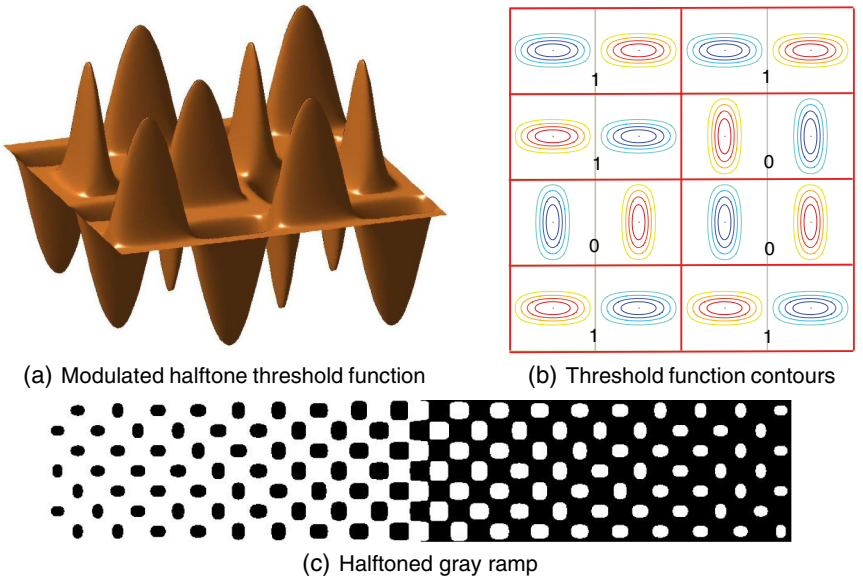

(a) Modulated halftone threshold function

(b) Threshold function contours

(c) Halftoned gray ramp

Fig. 12 Modulation of the halftone threshold function for data embedding. (a) Modulated halftone threshold function (4×2 halftone cell region), (b) contours for the modulated threshold function in (a), and (c) output halftones for a gray ramp image spanning the range from white to black. Observe the elliptical black (white) dots on white (black) backgrounds in highlight (shadows). Also note that at the 50% coverage mid graylevel the dots form a checkerboard with no discernible impact of the modulation.

hashes and localized embedding as has been done with cryptographic signatures for digital image authentication.[18,19] Such extensions are, however, nontrivial in our assessment. As already discussed in Sec. 1, robust hashes based on localized feature points in the image which have been proposed for digital image authentication in the presence of benign image modifications, tend to be particularly susceptible to print scan distortions. Robust hashes computed from less localized statistics, such as the scheme[27] that we benchmarked in Sec. 4, do not suffer from this limitation, however, their localization capability is consequently also limited. Attempts to make these latter methods more localized by computing signatures on smaller blocks reduce robustness and also have the potential for introducing vulnerability to the so-called vector-quantization attacks[40,41] and cryptographic weakness, if the computed robust hash is too short. Last, the scheme we propose in this paper also provides indication of the original authentic content in manipulated regions, something that is not feasible with robust hash based approaches.

For the demonstration of our proposed scheme, we used our high capacity data hiding method based on orientation modulation in clustered-dot halftones. As we noted in Sec. 3, alternative high capacity data embedding schemes could also be utilized within the our proposed framework. The capacity of the embedding scheme has significant implications for practical implementation and deployment in existing applications. Our demonstration in Sec. 4 used 8×8 in^2 printed images. In typical security applications, the size of the printed images is significantly smaller. We estimate that with our current data hiding scheme, a 2×2 in^2 offers adequate capacity for carrying an acceptable thumbnail of a facial portrait image, which is reasonable, although still larger than typical portrait images in passport and other security documents. Though beyond the scope of this paper, we note that enhancements to the high capacity data hiding schemes can be expected to bridge the capacity gap between our prototype demonstration and the requirements for practical deployment. Specifically, higher print resolution commonly available in security printing applications also enables higher capacity. For our orientation modulation scheme, with other parameters fixed, the capacity scales the square of the printer resolution in linear units (e.g., dpi) and, therefore, even a modest increase in resolution offers a significant capacity increase. The improvement in print resolution also increases capacity by enabling a greater degrees of freedom in the embedding, for instance by using four or more orientations for orientation modulation instead of the two in our scheme. Finally, color printing and capture can also boost capacity as shown in the recent work with barcodes.[42] It is also worth noting that because printed images are designed for human viewing, the "visual information content" in printed images and therefore the data size of an acceptable thumbnail scales in proportion to the printed area and is consequently well matched with the scaling of the capacity with the print area for a number of high-capacity print data hiding methods.

Our proposed method advances the state-of-the-art in printed document security. However, many open problems remain in this area. Specifically, we note that our method provides authentication of printed content and does not address applications requiring copy prevention or detection. Nor do we address the problem of embedding robust watermarks or fingerprints in printed documents, which are of interest in copyright enforcement. These and other areas remain challenging open fields for research.

Appendix A: Data Hiding, Extraction, and Image Registration

In this appendix, we provide a brief description (Additional details can be found in Ref. 30) of the specific high capacity data hiding scheme that we use for implementing our content authentication framework, emphasizing in particular, the method by which we achieve synchronization, which also serves to register the printed image to the thumbnail in the signature which is essential for the automated verification step in Sec. 3.3.

Figure 11 shows the overall block diagram of our signature transport scheme. The input message data, consisting of the encrypted thumbnail to be transported in the print, is represented by a binary vector $\mathbf{m}$. To achieve error-free communication of this message data in spite of the inevitable errors introduced in the embedding and extraction process, error control coding is employed. In our specific implementation, we use RA codes which are one of the current state-of-the-art class of codes that offer excellent performance and a flexible and simple design procedure accommodating a variety of rates. The RA encoder adds redundancy to the binary message data $\mathbf{m}$ to generate an encoded binary data vector $\mathbf{c}$. As the contone image $I(x, y)$ is halftoned to obtain a bilevel representation $I^h(x, y)$ suitable for printing, the coded data $\mathbf{c}$ is embedded in the halftone image. A hardcopy image is obtained by rendering the halftone image $I^h(x, y)$ on a printer, which can be utilized in the same manner as other prints that do not necessarily have data embedded. When authenticity of the print needs to be verified, the print is scanned and processed to recover the embedded message data; first, global and local synchronization are performed to overcome geometric distortions introduced in the print-scan processes, then bit-wise detection statistics corresponding to the embedded coded data $\mathbf{c}$ are calculated, which are then used in a belief propagation decoding algorithm for the RA code[43,44] to recover the embedded message. We note that like other "capacity-achieving" families of codes, the RA codes exhibit a threshold behavior. Once the effective signal-to-noise ratio of the print-scan channel exceeds a threshold, the data $\mathbf{m}$ is recovered without errors from the scanned image. Because $\mathbf{m}$ represents encrypted data for the thumbnail which would be rendered useless by even a few bit errors, the error-free recovery is crucial.

The data hiding technique that we employ performs the data embedding and halftoning jointly and, as is appropriate for a data-hiding method, emphasizes preservation of visual quality of the printed halftones over detectability of the embedded data. Specifically, the printed image is obtained by screening, i.e., point-by-point comparison of the contone image against a halftone threshold function. The halftone threshold function is modulated within each halftone cell by selecting, based on a bit of the encoded message $\mathbf{c}$, a threshold function that generates either horizontally or vertically oriented halftone dots. Figure 12(a) illustrates the halftone threshold function over a region corresponding to 4×2 halftone cells for random bits where the contours for the threshold function are shown in Fig. 12(b) with the corresponding modulation data super-imposed. From the

figure, one can see that continuity is maintained everywhere for the halftone threshold function. The actual halftone printed, and recoverability of individual bits, depend on the image content in addition to the halftone threshold function. The impact on the threshold function modulation and the image graylevel on the halftone dots is illustrated in Fig. 12(b) where a gray ramp image that goes smoothly from white to black along with the x-direction, which has been halftoned with a modulated threshold function, is shown in enlarged view showing the halftone dot structure at different graylevels. Black elliptical dots along with the two orientations are produced on a predominantly white background in highlights, i.e., regions of the contone image with less than 50% halftone area coverage, white elliptical "holes" with the two possible orientations are produced in the shadows, i.e., regions of the contone image with more than 50% halftone area coverage, and the 50% graylevel yields a checkerboard pattern where the modulation of the threshold does not impact the shapes of the dots. The joint halftoning and data embedding process generates a binary halftone image $I^h(x,y)$, which is then printed.

In addition to the impact of the image graylevel and various sources of noise introduced in the printing process, the recoverability of the embedded data from the print is also impacted by geometric distortions in the print-scan process, in particular global rotation and scaling introduced in the scanning process and locally varying geometric distortion that is typical in printing devices.[45,46] Parameters defining the global rotation and scaling are readily estimated by locating the peaks in the Fourier transform of the scanned image corresponding to the inherent 2-D (quasi) periodicity of the halftone dot structures in the printed image,[30,47] which then allows for compensation of the rotation and scaling. The local geometric distortion introduced by the printer varies spatially over the page in a smooth fashion which makes it imperceptible to viewers but is extremely significant at the scale of the halftone dot structures used for our data embedding, shifts of over half a cell over regions of the page are not atypical, which would clearly be catastrophic if ignored. We compensate for the local geometric distortion in conjunction with the data demodulation process,[30,48] in a manner analogous to decision directed synchronization in digital communications.[49] The local synchronization is particularly useful in the subsequent authentication of the printed content, we, therefore, summarize the demodulation and local synchronization process next.

To estimate the orientation of the halftone dot, we compute image moments of the scanned data $I^s(x,y)$ within the cell along with the two orthogonal embedding dimensions. Specifically, the image moment σ_x along with the X axis is computed within a halftone cell (Our description readily assumes less than 50% area coverage. The case for greater than 50% area coverage is readily handled by inverting the image) C as

$$\sigma_x = \frac{\sum_{x,y\in C} I^s(x,y)(x-\bar{x})^2}{\sum\limits_{x,y\in C} I^s(x,y)}, \tag{8}$$

where $\bar{x} = \sum_{x,y\in C} I^s(x,y)x / [\sum_{x,y\in C} I^s(x,y)]$ represents the X-coordinate of the center of mass of the halftone dot. The image moment σ_y along with the orthogonal y axis is calculated in a similar fashion. All pixels in the halftone cell contribute to image moments along with the horizontal and vertical directions according to their value $I^s(x,y)$ and distance to center of mass in the horizontal or vertical direction. The knowledge of the center of mass of the just demodulated halftone dot is then used to determine the position of the next halftone cell in the scan, thereby compensating for the local geometric distortions introduced in the printing process.[30,48] This synchronization process also automatically registers the printed image to the embedded thumbnail for the automated verification step in our proposed content authentication scheme, which is described in Sec. 3.3. Once image moments are calculated for all halftone cells in the image, error correction decoder estimates the message data $\hat{\mathbf{m}}$ using the calculated image moments. (Note that error correction decoder can also utilize the statistical model of the print-scan channel in the recovery of the message data. See Ref. 30 for details)

Acknowledgments

We would like to thank Dr. Vishal Monga who collaborated with us in our early work in exploring applications of high capacity data embedding in printed images.[31] We would also like to thank the anonymous reviewers for the manuscript for their comments and suggestions, which have significantly improved the presentation in the paper. This work was supported in part by a grant from New York State Office of Science, Technology & Academic Research (NYSTAR) through the Center for Electronic Imaging Systems (CEIS).

References

1. W. Kaule and G. Stenzel, "Security paper with authenticity features in the form of substances luminescing only in the invisible region of the optical spectrum, and process for testing the same," United States Patent No. 4,451,521 (1984).
2. M. Boehm, "Security paper," United States Patent No. 4,897,300 (1990).
3. R. D. Hersch, P. Emmel, and F. Collaud, "Reproduction of security documents, and color images with metallic inks," United States Patent No. 7,491,424 (2009).
4. J. C. Handley et al., "Magnetic watermark for text documents," United States Patent No. 7,386,159 (2008).
5. R. L. van Renesse, "Hidden and scrambled images: a review," *Proc. SPIE* **4677**, 333–348 (2002).
6. I. Amidror, S. Chosson, and R. D. Hersch, "Moiré methods for the protection of documents and products: a short survey," *J. Phys. Conf. Ser.* **77**(1), 012001 (2007).
7. J. Constant, "Holographic identification system using incoherent light," United States Patent No. 4,820,006 (1989).
8. S. Huang and J. K. Wu, "Optical watermarking for printed document authentication," *IEEE Trans. Inf. Foren. Sec.* **2**(2), 164–173 (2007).
9. B. Oztan and G. Sharma, "Continuous phase-modulated halftones," *IEEE Trans. Image Process.* **18**(12), 2718–2734 (2009).
10. B. Oztan and G. Sharma, "Per-separation clustered-dot color halftone watermarks: separation estimation based on spatial frequency content," *J. Electron. Imag.* **19**(4), 043007 (2010).
11. E. Métois et al., "FiberFingerprint identification," in *Proc. 3rd Workshop on Automatic Identification*, Tarrytown, New York, pp. 147–154 (2002).
12. B. Zhu, J. Wu, and M. S. Kankanhalli, "Print signatures for document authentication," in *Proc. 10th ACM Conf. Computer and Communications Security*, Washington, DC, pp. 145–154 (2003).
13. J. Picard, C. Vielhauer, and N. Thorwirth, "Towards fraud-proof ID documents using multiple data hiding technologies and biometrics," *Proc. SPIE* **5306**, 416–427 (2004).
14. S. J. Simske et al., "Spectral pre-compensation and security deterrent authentication," in *Proc. IS&T's NIP 24: Int. Conf. Digital Printing Technologies*, Pittsburgh, Pennsylvania, pp. 792–795 (2008).
15. L. O'Gorman and I. Rabinovich, "Secure identification documents via pattern recognition and public-key cryptography," *IEEE Trans. Pattern Anal. Mach. Intell.* **20**(10), 1097–1102 (1998).
16. K. Sengupta and D. MacIntyre, "The moment Mossad agents got their man?," The Independent [online] (17 February 2010).
17. A. Menezes, P. van Oorschot, and S. Vanstone, *Handbook of Applied Cryptography*, pp. 15–21, 25–30, 321–383, 559–561, CRC Press, Florida (1997).

18. P. W. Wong and N. Memon, "Secret and public key image watermarking schemes for image authentication and ownership verification," *IEEE Trans. Image Process.* **10**, 1593–1601 (2001).
19. M. U. Celik, G. Sharma, and A. M. Tekalp, "Lossless watermarking for image authentication: a new framework and an implementation," *IEEE Trans. Image Process.* **15**(4), 1042–1049 (2006).
20. S. Bhattacharjee and M. Kutter, "Compression-tolerant image authentication," in *Proc. IEEE Intl. Conf. Image Proc.*, Chicago, Illinois, pp. 435–439 (1998).
21. J. Fridrich and M. Goljan, "Robust hash functions for digital watermarking," in *Proc. IEEE Int. Conf. Information Technology: Coding and Computing (ITCC)*, Las Vegas, Nevada, pp. 173–178 (2000).
22. C. S. Lu and H. Y. M. Liao, "Structural digital signature for image authentication: an incidental distortion resistant scheme," *IEEE Trans. Multimed.* **5**(2), 161–173 (2003).
23. V. Monga and B. Evans, "Perceptual image hashing via feature points: performance evaluation and tradeoffs," *IEEE Trans. Image Process.* **15**(11), 3452–3465 (2006).
24. P. Roetling and R. P. Loce, "Digital halftoning," Chapter 10 in *Image Processing: Fundamentals and Applications*, E. R. Dougherty, Ed., Marcel Dekker, New York (1994).
25. R. Rauch and M. Rahnavard, "Raster output scanning system having scan line non-linearity compensation means," United States Patent No. 6,178,031 (2001).
26. J. Appel, "ROS bow compensation," United States Patent No. 6,232,991 (2001).
27. J. Barr, B. Bradley, and B. T. Hannigan, "Using digital watermarks with image signatures to mitigate the threat of the copy attack," in *Proc. IEEE Int. Conf. Acoustics Speech and Sig. Proc.*, Hong Kong, Vol. 3, III–69 (2003).
28. V. Monga, D. Vats, and B. Evans, "Image authentication under geometric attacks via structure matching," in *IEEE Int. Conf. Multimedia and Expo*, pp. 229–232, IEEE, Amsterdam, The Netherlands (2005).
29. L. Yu and S. Sun, "Image authentication in print-and-scan scenario," in *Int. Conf. Intelligent Information Hiding and Multimedia Signal Processing*, pp. 295–298, IEEE Computer Society, Kaohsiung, Taiwan (2007).
30. O. Bulan, G. Sharma, and V. Monga, "Orientation modulation for data hiding in clustered-dot halftone prints," *IEEE Trans. Image Process.* **19**, 2070–2084 (2010).
31. O. Bulan, G. Sharma, and V. Monga, "Application of high capacity data hiding in halftone images," in *Proc. IS&T's NIP 24: Int. Conf. Digital Printing Technologies*, Pittsburgh, Pennsylvania, pp. 787–791 (2008).
32. M. Kutter, S. V. Voloshynovskiy, and A. Herrigel, "Watermark copy attack," *Proc. SPIE* **3971**, 371–380 (2000).
33. J. Nechvatal et al., "Report on the development of the Advanced Encryption Standard (AES)," *J. Res. Natl. Inst. Stand. Technol.* **106**, 511–576 (2001).
34. D. S. Taubman and M. W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*, Kluwer Academic Publishers, Boston (2002).
35. M. Rabbani and R. Joshi, "An overview of the JPEG2000 still image compression standard," *Signal Proc. Image Commun.* **17**(1), 3–48 (2002).
36. N. Damera-Venkata et al., "Hardcopy image barcodes via block-error diffusion," *IEEE Trans. Image Process.* **14**, 1977–1989 (2005).
37. R. Ulichney, M. Gaubatz, and S. J. Simske, "Encoding information in clustered-dot halftones," in *Proc. IS&T's NIP 26: Int. Conf. Digital Printing Technologies*, Austin, Texas, pp. 602–605 (2010).
38. S. C. Pei and J. M. Guo, "High-capacity data hiding in halftone images using minimal-error bit searching and least-mean square filter," *IEEE Trans. Image Process.* **15**(6), 1665–1679 (2006).
39. Y.-Y. Chen et al., "Stegatone performance characterization," *Proc. SPIE* **8665**, 86650Q–86650Q (2013).
40. M. Holliman and N. Memon, "Counterfeiting attacks on oblivious block-wise independent invisible watermarking schemes," *IEEE Trans. Image Process.* **9**(3), 432–441 (2000).
41. M. U. Celik et al., "Hierarchical watermarking for secure image authentication with localization," *IEEE Trans. Image Process.* **11**(6), 585–595 (2002).
42. H. Blasinski, O. Bulan, and G. Sharma, "Color barcodes for mobile applications: a per channel framework," *IEEE Trans. Image Process.* **22**, 1498–1511 (2013).
43. D. Divsalar, H. Jin, and R. J. McEliece, "Coding theorems for "turbo-like" codes," in *Proc. Allerton Conf.*, pp. 201–210, IEEE, Monticello, Illinois (1998).
44. H. Jin, A. Khandekar, and R. McEliece, "Irregular repeat-accumulate codes," in *Proc. 2nd Int. Symp. on Turbo Codes and Related Topics*, pp. 1–8, Ecole Nationale Superieure des Telecommunications de Bretagne, Brest, France (2000).
45. T. Kazama and Y. Matsuzaki, "Image forming apparatus, image forming method, and storage medium storing program for forming image," United States Patent Application Publication No. US2007/0139715 A1 (2007).
46. D. N. Curry, "Two dimensional linearity and registration error correction in a hyperacurity printer," United States Patent No. 5,732,162 (1998).
47. K. Solanki et al., "'Print and scan' resilient data hiding in images," *IEEE Trans. Inf. Foren. Sec.* **1**, 464–478 (2006).
48. O. Bulan and G. Sharma, "High capacity color barcodes: per channel data encoding via orientation modulation in elliptical dot arrays," *IEEE Trans. Image Process.* **20**(5), 1337–1350 (2011).
49. J. G. Proakis, *Digital Communications*, 4th ed., McGraw-Hill, New York (2001).

Orhan Bulan received a BS degree with high honors in electrical and electronics engineering from Bilkent University, Ankara, Turkey, in 2006, and MS and PhD degrees in electrical and computer engineering from University of Rochester, NY, in 2007 and 2012, respectively. He is currently a post-doctoral fellow in the Xerox Research Center Webster, Webster, NY. He was with Xerox during the summers of 2009, 2010, and 2011 as a research intern. He is the recipient of the best student paper award at the 2008 Western New York Image Processing Workshop organized by the Rochester Chapter of the IEEE Signal Processing Society. His recent research interests include signal/image processing, video processing, computer vision, and machine learning. He has four issued patents and over 10 pending patent applications in these areas.

Gaurav Sharma is an associate professor at the University of Rochester in the Department of Electrical and Computer Engineering, Department of Biostatistics and Computational Biology, and Department of Oncology. From 2008 to 2010, he served as the director for the Center for Emerging and Innovative Sciences (CEIS), a New York state funded center for promoting joint university-industry research and technology development, which is housed at the University of Rochester. From August 1996 to August 2003, he was with Xerox Research and Technology, in Webster, NewYork, initially as a member of research staff and subsequently at the position of principal scientist. Dr. Sharma's research interests include color science and imaging, multimedia security and watermarking, and bioinformatics. He is the editor of the "*Color Imaging Handbook*," published by CRC Press in 2003. He is a fellow of SPIE, of the Society of Imaging Science and Technology (IS&T) and of IEEE, and a member of Sigma Xi, Phi Kappa Phi, Pi Mu Epsilon honor societies. He served as thesymposium chair for the 2012 SPIE/IS&T Electronic Imaging Symposium and as a technical program chair for the 2012 IEEE International Conference on Image Processing (ICIP). He is the editor-in-chief for the *Journal of Electronic Imaging* and in the past has served as an associate editor for the *Journal of Electronic Imaging, IEEE Transactions on Image Processing,* and *IEEE Transactions on Information Forensics and Security*.