COMPRESSIVE SENSING-BASED IMAGE HASHING+

Li-Wei Kang, *Chun-Shien Lu*,* *and Chao-Yung Hsu*

Institute of Information Science, Academia Sinica, Taipei, Taiwan, ROC

ABSTRACT

In this paper, a new image hashing scheme satisfying robustness and security is proposed. We exploit the property of dimensionality reduction inherent in compressive sensing/sampling (CS) for image hash design. The gained benefits include (1) the hash size can be kept small and (2) the CS-based hash is computationally secure. We study the use of visual information fidelity (VIF) for hash comparison under Stirmark attacks. We further derive the relationships between the hash of an image and both of its MSE distortion and visual quality measured by VIF, respectively. Hence, based on hash comparisons, both the distortion and visual quality of a query image can be approximately estimated without accessing its original version. We also derive the minimum distortion for manipulating an image to be unauthentic to measure the security of our scheme.

Index Terms—Authentication, Compressive sensing, Image hashing, Robustness, Security, and Visual information fidelity.

1. INTRODUCTION

A media hash is a content-based compact representation of a digital multimedia signal, which has been applicable extensively. Usually, an image hashing scheme should meet the following five requirements [1]-[5]: (1) perceptual robustness; (2) collision resistance; (3) one-way and unpredictability; (4) database retrieval efficiency; and (5) manipulated image quality estimation. Recently, compressive sensing (CS) [6]-[12] has been successfully applicable to image authentication [4]-[5] based on its superiority in random projection for sparse signal and computational security. In [4], an image authentication scheme based on CS and distributed source coding (DSC) was proposed, where the image hash is derived from the DSC-encoded quantized random projection coefficients of an image. To perform authentication, a DSC decoder decodes the received hash bits with the test image serving as the side information, where the authenticity depends on the success/fail of the DSC decoding. A similar scheme was also proposed in [5].

In this paper, a robust and secure image hashing scheme via CS and visual information fidelity (VIF) is proposed, which can meet the above-mentioned five requirements. The novelties and advantages from exploiting CS and VIF of our scheme can be summarized as follows: (1) low-complexity hash extraction: the hash can be simultaneously extracted while acquiring an image using the CS single-pixel imaging camera [9] and computational security can be guaranteed without performing additional randomization process; (2) short hash length: coming from the highly sparse signal dimensionality reduction capability of CS; (3) advanced hash comparison: the highly sparse signal reconstruction capability for hash comparison; (4) perceptual robustness: the hash comparison via VIF is robust against attacks; and (5) the estimations for both image distortion and visual quality via hash comparisons. The major differences distinguishing our scheme from [4]-[5] include: (1) both the random projection and reconstruction capabilities for sparse signal of CS are exploited in our scheme, whereas only random projection is exploited in [4]-[5]; (2) in terms of hash comparison, both the two metrics, MSE and VIF, are exploited in our scheme, whereas DSC decoding is exploited in [4]-[5]; and (3) both image distortion and visual quality can be estimated in our scheme, whereas only distortion can be estimated in [5].

2. BACKGROUND

2.1. Compressive sensing (CS)

Assume that a basis matrix Ψ with size $N\times N$ can provide a K-sparse representation for a real value signal (*e.g.*, image signal) x with length N. That is, $x = \Psi\theta$ and θ with length N can be well approximated using only $K << N$ non-zero entries. CS [6] states that x can be accurately reconstructed by taking only $M = O(K\log(N/K))$, $K < M << N$, linear measurements from:

$$y = \Phi x, \tag{1}$$

where y is an $M\times 1$ measurement vector, and Φ is an $M\times N$ measurement matrix controlled by a secret key, which is incoherent with Ψ. The M measurements in y can be viewed as the compressed and encrypted version of x. We have found that CS, achieving dimensionality reduction, is useful to represent a hash.

2.2. Visual information fidelity (VIF)

In [13], a novel image quality assessment, called visual information fidelity (VIF), ranged from 0 to 1, is proposed. To quantify the visual quality of a distorted image, an image information measure is proposed to quantify the information presented in its original version and how much of this information can be extracted from the distorted image. In particular, VIF models images in the wavelet domain using Gaussian scale mixtures, where a scale-space-orientation wavelet decomposition, called the steerable pyramid [14], is used, which has been shown to be translation- and rotation-invariant. We have found that these properties are suitable to measure the similarity between a geometric-manipulated image and its original version.

3. PROPOSED IMAGE HASHING SCHEME

3.1. Hash vector extraction via CS

Our image hashing scheme is shown in Fig. 1. First, a pre-processing step is conducted by first converting an image x of size $n\times n$ to be gray level, followed by down-sampling with ratio $1/B^2$ to the image x_o of size $N = (n/B)\times(n/B)$. For hash extraction, given an $M\times N$ measurement matrix Φ, $M << N$, controlled by a secret key S, x_o is randomly projected (Eq. (1)) to a measurement vector with size M, where each hash value (vector component) is quantized to form the final "hash vector" $y = [y_1, y_2, \ldots, y_M]^T$, followed by entropy-encoding. For a received image x' to be authenticated, the same hash extraction process with the same key is applied to generate the hash vector $y' = [y'_1, y'_2, \ldots, y'_M]^T$. In this paper, the exploited measurement matrix Φ is the scrambled block Hadamard ensemble (SBHE) matrix [10]. The basis matrix Ψ used here is DWT (discrete wavelet transform) basis. For hash value quantization and encoding, a non-uniform quantizer (designed by the k-means algorithm) and a Huffman table are created.

*Corresponding author: lcs@iis.sinica.edu.tw

+This work was supported in part by National Science Council, Taiwan, ROC, under Grant NSC 97-2628-E-001-011-MY3.

3.2. Hash vector comparison and distortion estimation

To compare two image hash vectors, $y = [y_i]^T$ and $y' = [y'_i]^T$, $i = 1, 2, \ldots, M$, we simply calculate the MSE between them via

$$\mathrm{MSE}(y', y) = (1/M)\sum_{i=1}^{M}(y'_i - y_i)^2. \quad (2)$$

To estimate the distortion between x' and x, *i.e.*, MSE(x', x), from MSE(y', y), which has also been mentioned in [5], we derive the relationship between them for our hash scheme as follows, and verify it via simulations in Sec. 4. Consider the down-sampled versions, $x'_o = \Psi\theta'_o$ and $x_o = \Psi\theta_o$, of x' and x, respectively, where Ψ is the DWT basis, $e_o = \theta_o' - \theta_o$, $y' = \Phi x'_o$, $y = \Phi x_o$, where Φ is the SBHE measurement matrix, and $A = \Phi\Psi$. For simplicity, without taking the quantization of measurements into account, we have:

$$\|y' - y\|_2^2 = \|\Phi(x'_o - x_o)\|_2^2 = \|\Phi\Psi(\theta'_o - \theta_o)\|_2^2 = \|Ae_o\|_2^2.$$

Based on the assumption that A obeys the restricted isometry property (RIP) [6], we have

$$(1-\delta_k)\|e_o\|_2^2 \le \|Ae_o\|_2^2 \le (1+\delta_k)\|e_o\|_2^2,$$

where δ_k is the isometry constant of A for all k-sparse vectors e_o, and δ_k is not too close to 1. Hence, we have

$$\|y' - y\|_2^2 = \|Ae_o\|_2^2 \approx \|e_o\|_2^2 = \|\theta'_o - \theta_o\|_2^2 \approx \|x'_o - x_o\|_2^2,$$

and $\mathrm{MSE}(x_o', x_o) \approx \mathrm{MSE}(y', y)$. (3)

We observe the relationship between MSE(x', x) and MSE(x_o', x_o) via the correlation coefficient between them, which is usually large enough. Hence, by applying linear regression technique, we have

$$\mathrm{MSE}(x', x) \approx \alpha_0 + \alpha_1 \mathrm{MSE}(x'_o, x_o) + \alpha_2[\mathrm{MSE}(x'_o, x_o)]^2 + \cdots + \alpha_t[\mathrm{MSE}(x'_o, x_o)]^t, \quad (4)$$

where α_0, α_1, …, α_t can be estimated via least squares estimation using training images.

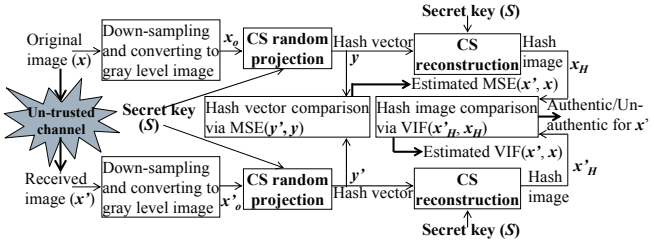

Fig. 1. Our image hashing scheme.

Based on RIP and Eqs. (3)-(4), we can claim that comparing the MSE between two hash vectors can be approximately equivalent to comparing the MSE of the two corresponding images. To decide whether x' is manipulated from x based on MSE(y', y), intuitively, if MSE(y', y) $\le \tau_{mse}$, where τ_{mse} is a predefined threshold, x' can be authentic; otherwise, x' is unauthentic. Although, simply calculating the MSE between two hash vectors doesn't take geometric manipulations into account, we will show that the authentication process described as follows can be robust to most non-geometric and slight geometric manipulations (Sec. 4) and provide a guideline for our advanced authentication process described in Sec. 3.3. By considering the two hash vectors, y and y', respectively, extracted from x and x', and letting v_{mse} = MSE(y', y), we model the image authentication problem as a hypothesis testing problem with two hypotheses: H_0: x' is authentic, *i.e.*, $v_{mse} \le \tau_{mse}$, and H_1: x' is unauthentic, *i.e.*, $v_{mse} > \tau_{mse}$. Because each CS measurement is a random projection of an image, it can be approximated with i.i.d. Gaussian distribution [12]. Based on central limit theorem, MSE(y', y) can be well approximated by a Gaussian distribution when $M \ge 30$. Hence, v_{mse} = MSE(y', y) can be viewed as a random variable with i.i.d Gaussian distribution. Then, the probability density functions, $f(v_{mse}|H_0)$ and $f(v_{mse}|H_1)$, can be, respectively, expressed as:

$$f(v_{mse} \mid H_0) = \left(1/\sigma_{mse\_1}\sqrt{2\pi}\right)e^{-(v_{mse}-\mu_{mse\_1})^2/2\sigma_{mse\_1}^2}, \quad (5)$$

$$f(v_{mse} \mid H_1) = \left(1/\sigma_{mse\_2}\sqrt{2\pi}\right)e^{-(v_{mse}-\mu_{mse\_2})^2/2\sigma_{mse\_2}^2}, \quad (6)$$

where μ_{mse_1}, σ_{mse_1}, and μ_{mse_2}, σ_{mse_2}, respectively, denote the means and standard deviations of v_{mse} when x' and x are relevant, and when x' and x are irrelevant. Hence, the two hypotheses can be re-expressed as:

$$H_0: f(v_{mse}|H_0) \ge f(v_{mse}|H_1), \text{ and } H_1: f(v_{mse}|H_0) < f(v_{mse}|H_1). \quad (7)$$

By substituting Eqs. (5)-(6) into H_0 in Eq. (7), we can obtain that when x' is authentic, the optimal v_{mse} can be calculated as:

$$v_{mse\_opt} = \left(b \pm \sqrt{b^2 - ac}\right)/a, \quad (8)$$

where $a = \left(\sigma_{mse\_1}^2 - \sigma_{mse\_2}^2\right)$, $b = \left(\sigma_{mse\_1}^2\mu_{mse\_2} - \sigma_{mse\_2}^2\mu_{mse\_1}\right)$, and

$$c = \left(\sigma_{mse\_1}^2\mu_{mse\_2}^2 - \sigma_{mse\_2}^2\mu_{mse\_1}^2 - 2\sigma_{mse\_1}^2\sigma_{mse\_2}^2\ln\left(\sigma_{mse\_1}/\sigma_{mse\_2}\right)\right).$$

Hence, τ_{mse} can be set to $v_{mse\_opt} > 0$. The main goal of our hash vector comparison via MSE in the first stage of our scheme is to estimate the distortion of the image to be authenticated, and provide a benchmark for advanced hash comparison (hash image comparison) in the second stage, as described in Sec. 3.3.

3.3. Hash image comparison and visual quality estimation

To resist geometric manipulations, we propose to compare the CS reconstructed images (called hash images) via VIF. With the benefit of the translation-invariance and rotation-invariance of the steerable pyramid wavelet decomposition [14] employed in VIF, we have observed that VIF can exhibit higher similarity between an image and its geometric-manipulated versions than that between it and other irrelevant ones. Here, a hash image is generated from its hash vector using the gradient projection for sparse reconstruction algorithm [7], which solves the convex unconstrained optimization problem for CS reconstruction. Let x'_H and x_H be, respectively, the hash images of the images x' and x, and let v_{vif} = VIF(x'_H, x_H). If $v_{vif} \ge \tau_{vif}$, where τ_{vif} is a predefined threshold, x' can be authentic; otherwise, x' is unauthentic. Different from MSE, the larger v_{vif} is, the more similar x' and x are. The threshold τ_{vif} can also be similarly derived via hypothesis testing by assuming v_{vif} to be a random variable with i.i.d. Gaussian.

On the other hand, we have observed that the correlation coefficient for VIF(x'_H, x_H) and VIF(x', x) is usually large enough. That is, the hash of an image can estimate the visual quality (measured by VIF) of the image itself without needing to access its original version. By applying linear regression technique, VIF(x', x) can be estimated from VIF(x'_H, x_H) via

$$\mathrm{VIF}(x', x) \approx \beta_0 + \beta_1 \mathrm{VIF}(x'_H, x_H) + \beta_2[\mathrm{VIF}(x'_H, x_H)]^2 + \cdots + \beta_r[\mathrm{VIF}(x'_H, x_H)]^r, \quad (9)$$

where β_0, β_1, …, and β_r can be estimated via least squares estimation using training images.

3.4. Robustness and collision resistance of our scheme

The robustness and collision resistance of our hash scheme can be analyzed via quantifying the true positive rate (TPR) and false positive rate (FPR). The TPR $P_T(\tau_{vif})$ for our hash image comparison in Sec. 3.3 can be expressed as:

$$P_T(\tau_{vif}) = \Pr(v_{vif} \mid H_0 \ge \tau_{vif}) = 1 - \int_{z=-\infty}^{z=(\tau_{vif}-\mu_{hash\_1})/\sigma_{hash\_1}} \left(1/\sqrt{2\pi}\right)e^{-z^2/2}dz. \quad (10)$$

In addition, the FPR $P_F(\tau_{vif})$ can be expressed as:

$$P_F(\tau_{vif}) = \Pr(v_{vif} \mid H_1 \ge \tau_{vif}) = 1 - \int_{z=-\infty}^{z=(\tau_{vif}-\mu_{hash\_2})/\sigma_{hash\_2}} \left(1/\sqrt{2\pi}\right)e^{-z^2/2}dz, \quad (11)$$

where μ_{hash_1}, σ_{hash_1}, and μ_{hash_2}, σ_{hash_2}, respectively, denote the means and standard deviations of v_{vif} when x' and x are relevant,

and when x' and x are irrelevant. Similar derivations are also valid for our hash vector comparisons.

3.5. Security of our scheme

The security of our hash scheme is guaranteed by the inherent computational security in CS [8]. Without knowing the secret key, an attacker cannot get the correct measurement matrix, generate the valid hash vector for an image, and reconstruct the hash image from a hash vector. Even if an attacker knowing a key pool including the correct key wants to try all of them to reconstruct the correct hash image, it is still hard to decide which key can result in a successful reconstruction.

On the other hand, the differential entropy (DE) has been employed to evaluate the security of an image hashing scheme [1]-[2]. The DE for each hash value y_i (Gaussian distributed with mean μ_i and variance σ_i^2) in a hash vector $y = [y_1, y_2, \ldots, y_M]^T$ can be calculated as: $h(y_i) = (1/2)\log_2(2\pi e\sigma_i^2)$. By considering a whole hash vector y, its DE can be calculated as: $h(y) = (1/2)\log_2[(2\pi e)^M |Q|]$, where $|Q|$ is the determinant of the covariance matrix Q of y. Higher DE implies that the secret key cannot be easily estimated. The DE values obtained using our scheme will be discussed in Sec. 4. The above-mentioned properties can meet the one-way and unpredictability requirements of media hashing.

In addition, let's consider a scenario that an attacker may want to modify a copyrighted image so that its hash cannot be identified, but wish that its quality can be kept. Consider an original image x, its modified version x', their respective down-sampled versions with length N, x_o and x'_o, and their respective hash vectors with length M ($M << N$), $y = \Phi x_o$ and $y' = \Phi x'_o$, where Φ is the SBHE measurement matrix. An attacker should wish that MSE(x', x) is acceptable while MSE(y', y) > τ_{mse}, so that he/she can use x' without being detected. The required minimum MSE(x', x) when MSE(y', y) > τ_{mse} can be derived as follows. Based on Eq. (3) we have MSE(x_o', x_o) ≈ MSE(y', y). To simplify derivations by assuming MSE(x', x) ≈ ρMSE(x_o', x_o), where ρ is a coefficient, we have:

$$\text{MSE}(x', x) \approx \rho\text{MSE}(x_o', x_o) \approx \rho\text{MSE}(y', y) > \rho\tau_{mse}. \quad (12)$$

Hence, the minimum MSE(x', x) that x' is unauthentic is $\rho\tau_{mse}$. Based on our simulations from several attacks, if an image is decided to be unauthentic, the minimum MSE is usually MSE(x', x) > 500, and meanwhile VIF(x', x) < 0.02. Hence, even if an attacker can successfully modify a copyrighted image to be unauthentic, the modified image's quality is usually unacceptable.

4. SIMULATION RESULTS

To evaluate our image hashing scheme, we conducted four kinds of simulations, respectively, shown in Secs. 4.1-4.4. Ten 512×512 standard test images were used as the training images for estimating τ_{mse} and τ_{vif}, where the obtained optimal τ_{mse} and τ_{vif} are 1605.58 and 0.0917, respectively. For hash value quantization and encoding, 70 training images were used to generate a non-uniform codebook with 32 codewords, where each codeword is on average represented by 4.78 bits via Huffman codes. Here, each 512×512 (n = 512) color test image x was first converted to 256 gray-levels and down-sampled to 16×16 image x_o (B = 32 and N = 256). Then, the 77×256 (M = 77) SBHE matrix with the secret key S = 3587642 [10] was used to randomly project x_o via Eq. (1) to get the measurement vector, where each component is quantized to form the final hash vector $y = [y_1, y_2, \ldots, y_{77}]^T$, followed by entropy-encoding using the codebook. Hence, the average hash length for a 512×512 image is about 368.06 bits. The ten 512×512 standard test images (*e.g.*, *Baboon*, *F16*, *Lena*, *Pepper*, etc.) outside the above-mentioned training images were used to evaluate our scheme.

4.1. Robustness and collision resistance evaluation

The Stirmark 3.1 and 4.0 benchmarks [15], including total 203 geometric and non-geometric manipulations (*e.g.*, compression, brightness/contrast adjusting, noising, cropping, scaling, and rotation), were used to evaluate the robustness of our scheme. For evaluating the TPR, hash comparisons were conducted between each image and its 203 manipulated versions. For evaluating the FPR, we compare the hash for each image and the 203 manipulated versions of each of the other nine images. The real and theoretical TPR and FPR values for our hash vector comparison are 0.9064 and 0.0795 (based on the optimal τ_{mse}), and 0.9082 and 0.0778, respectively, which are accurately matched. The real and theoretical TPR and FPR values for our hash image comparison are 0.9094 and 0.0684 (based on the optimal τ_{vif}), and 0.9793 and 0.0537, respectively, which are somewhat different because VIF values may be not exactly i.i.d. Gaussian.

The ROC curves (TPR-FPR curves) [1]-[2] obtained from our two hash comparison strategies by adjusting the respective two thresholds, τ_{mse} and τ_{vif}, and the "feature points hash" scheme [2] for the ten test images are shown in Fig. 2. It can be observed that the performance of hash image comparison outperforms that of hash vector comparison. The main reason is that the hash vector comparison only calculates the MSE between two hash vectors without considering geometric image manipulations, whereas the hash image comparison calculates the VIF between two CS reconstructed images, which is robust to several geometric image manipulations. It can also be observed that the performances of our two hash schemes can significantly outperform the "feature points hash" scheme [2]. The authentication performances vs. each of the 203 Stirmark manipulations for the *Baboon* images (vs. *Lena* images) are shown in Figs. 3-4, where it is hoped that the black ("+") and pink ("o") curves can be well separated by the blue ("-") curve. It can be found from Figs. 3-4 that the "hash vector MSE" and "hash image VIF" between the *Baboon* image and its manipulated versions can be well discriminated from those between the *Baboon* image and the corresponding manipulated versions of the *Lena* image.

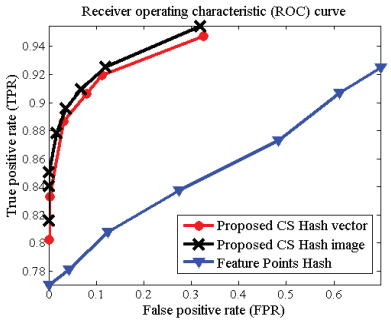

Fig. 2. ROC curves obtained using our scheme.

4.2. Distortion and visual quality estimation

The true and estimated MSE values, and the true and estimated VIF values, between *Baboon* image and each of its 203 manipulated versions are, respectively, shown in Figs. 3-4, where it is hoped that the green ("-.") curve can approximately estimate the red ("×") curve which is unavailable in the real situations. It can be observed from Figs. 3-4 that both the distortion measured by MSE and the visual quality measured by VIF of an image can be approximately estimated via our hash scheme for most image manipulations.

4.3. Evaluation of image database retrieval efficiency

We created a database consisting of 20000 images, including the 203 manipulated versions of each of the ten test images and several other images. We used each test image as a query image to find the top most similar 203 images, and evaluate the precision as *NC*/203, where *NC* denotes the number of retrieved images truly belonging to the manipulated versions of the query image. First, we roughly find the top *NV* images, $203 < NV << 20000$, with the smallest MSEs between their respective hash vectors and that of the query image. Then, we carefully evaluate the retrieved *NV* images by calculating their VIFs between their respective hash images and that of the query image to find the final top 203 images. The average precisions (%) for the ten images are listed in Table 1, where it can be observed that comparison using MSE only has been fairly good and approximate the overall precision. The main reasons include: (1) different from the authentication process, the comparison using MSE only wants to find the top *NV* images without caring about whether the MSEs are smaller than τ_{mse}; (2) some severely manipulated images may not be retrieved in the top *NV* images; and (3) VIF still cannot be robust to some severe manipulations (*e.g.*, severe noising, severe cropping, severe scaling, and severe rotation operations).

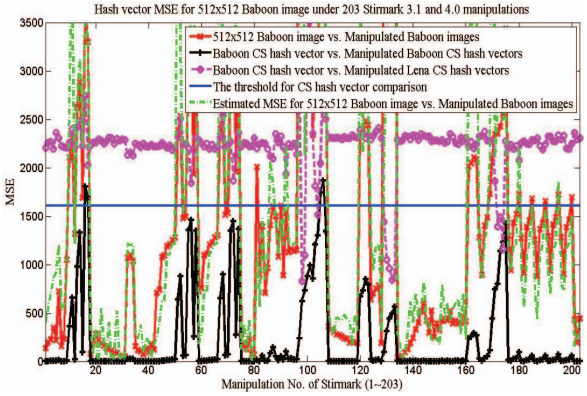

Fig. 3. The performances of hash vector comparison (MSE).

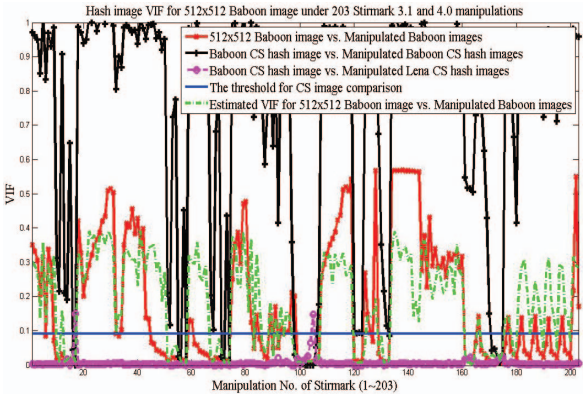

Fig. 4. The performances of hash image comparison (VIF).

4.4. Security evaluation

We randomly selected 5000 secret keys (*S*) to generate the hash vector, and calculate the average differential entropy (ADE) for each hash value and the differential entropy (DE) for each whole hash vector. The ADE and DE values for the three images are shown in Table 2. It can be found that the ADE values of our scheme are slightly smaller than or comparable to the ADE values reported in [1] (without considering hash value quantization). The main reason is that our hash is just only extracted by CS random projection, followed by quantization, whose entropy can be intentionally increased by performing an additional randomization process [2]. In addition, the minimum required distortion (MRD) measured by MSE for the three images being unauthentic and the respective corresponding VIFs are also shown in Table 2, which are obtained from our simulations. Please note that there is currently no direct relationship between MRD and VIF.

Table 1. Average image retrieval precisions.

NV	MSE only	Overall precision
250	87.19%	87.49%
500	87.19%	88.23%

Table 2. ADE, DE, and MRD values obtained using our scheme.

DE values	*Baboon*	*Lena*	*Pepper*
ADE	6.90	7.24	7.41
DE	772.22	798.44	811.49
MRD (VIF)	661.98 (0.0092)	592.95 (0.0152)	543.41 (0.0047)

5. CONCLUSIONS

In this paper, we exploit the characteristics of CS and VIF to propose a robust and secure image hashing scheme. With the benefit of CS, the hash size can be kept small, and the scheme is computationally secure. With the benefit of VIF, our scheme is robust to most image manipulations. We also derive the relationships between the hash of an image and both of its distortion and visual quality, respectively. On the other hand, we study the required minimum distortion for manipulating an image to be unauthentic to show the security of our scheme.

REFERENCES

[1] A. Swaminathan, Y. Mao, and M. Wu, "Robust and secure image hashing," *IEEE Trans. Inf. Forensics Security*, vol. 1, no. 2, June 2006.

[2] V. Monga and B. L. Evans, "Perceptual image hashing via feature points: performance evaluation and tradeoffs," *IEEE Trans. on Image Processing*, vol. 15, no. 11, pp. 3453-3466, Nov. 2006.

[3] P. J. O. Doets and R. L. Lagendijk, "Distortion estimation in compressed music using only audio fingerprints," *IEEE Trans. on Audio, Speech, and Language Processing*, vol. 16, no. 2, Feb. 2008.

[4] Y. C. Lin, D. Varodayan, and B. Girod, "Distributed source coding authentication of images with contrast and brightness adjustment and affine warping," *Proc. of Picture Coding Symposium*, USA, May 2009.

[5] M. Tagliasacchi, G. Valenzise, and S. Tubaro, "Hash-based identification of sparse image tampering," to appear in *IEEE Trans. on Image Processing*.

[6] E. J. Candes and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, 2008.

[7] M. A. T. Figueiredo et al., "Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems," *IEEE J. of Selected Topics in Signal Processing*, 2007.

[8] Y. Rachlin and D. Baron, "The secrecy of compressed sensing measurements," *Proc. Allerton*, IL, USA, Sept. 2008.

[9] M. F. Duarte et al., "Single-pixel imaging via compressive sampling," *IEEE Signal Processing Magazine*, vol. 25, pp. 83-91, Mar. 2008.

[10] L. Gan, T. T. Do, and T. D. Tran, "Fast compressive imaging using scrambled hadamard ensemble," *Proc. EUSIPCO*, Switzerland, 2008.

[11] L. W. Kang and C. S. Lu, "Distributed compressive video sensing," *Proc. IEEE ICASSP*, Taipei, Taiwan, ROC, Apr. 2009.

[12] V. Cevher et al., "Compressive sensing for background subtraction," *Proc. European Conf. on Computer Vision*, France, Oct. 2008.

[13] H. R. Sheikh and A. C. Bovik, "Image information and visual quality," *IEEE Trans. Image Process.*, vol. 15, no. 2, pp. 430-444, Feb. 2006.

[14] E. P. Simoncelli and W. T. Freeman, "The steerable pyramid: a flexible architecture for multi-scale derivative computation," *Proc. ICIP*, 1995.

[15] F. A. P. Petitcolas, "Watermarking schemes evaluation," *IEEE Signal Processing Magazine*, vol. 17, no. 5, pp. 58-64, Sept. 2000.