A Framework for Theoretical Analysis of Content Fingerprinting

Avinash L. Varna, Wei-Hong Chuang, and Min Wu

Department of Electrical and Computer Engineering and
Institute for Advanced Computer Studies,
University of Maryland, College Park.
{varna, whchuang, minwu}@umd.edu

ABSTRACT

The popularity of video sharing platforms such as Youtube has prompted the need for the development of efficient techniques for multimedia identification. Content fingerprinting is a promising solution for this problem, whereby a short "fingerprint" that captures robust and unique characteristics of a signal is computed from each multimedia document. This fingerprint is then compared with a database to identify the multimedia. Several fingerprinting techniques have been proposed in the literature and have been evaluated using experiments. To complement these experimental evaluations and gain a deeper understanding, this paper proposes a framework for theoretical modeling and analysis of content fingerprinting schemes. Analysis of some key modules for fingerprint encoding and matching are also presented under this framework.

Keywords: Content fingerprinting, theoretical modeling.

1. INTRODUCTION

Online services such as Flickr and YouTube have made content sharing popular. However, as users may upload copyrighted content to these websites, concerns about illegal content modification, duplication and digital rights infringement have also been raised. Content fingerprinting is a promising solution to this problem. Fingerprints are compact signatures that capture robust and unique features of multimedia that can be used for efficient identification. Several content fingerprinting schemes that capture various aspects of multimedia signals robust to different kinds of processing have been proposed in the literature, a survey of which may be found in [1].

Traditionally, these fingerprinting schemes have been evaluated using experiments on moderate-sized benchmark databases. From these experimental evaluations, it is difficult to infer how the performance of these fingerprinting schemes would scale to large databases with billions of video. Theoretical modeling and analysis can complement these experimental evaluations and help us in developing a deeper understanding of fingerprinting schemes. Such an analysis can also guide the design of better fingerprinting schemes. A systematic study would also help identify weaknesses of fingerprinting systems that may be exploited by smart attackers to circumvent the system and allow suitable counter-attack strategies to be devised.

While theoretical studies have been previously carried out in the related field of robust image hashing [2, 3], there are subtle differences in content fingerprinting that bring with them unique challenges. Traditionally, robust hashing was studied in the context of authentication, where the main objective was to prevent an adversary from forging an image that has the same hash as a given image or video. In contrast, while collisions or false alarms are also a concern in content fingerprinting, the main threat model is an adversary making minor modifications to a given multimedia document that would result in a significantly different fingerprint and prevent identification. Another difference between fingerprinting and robust hashing is that fingerprinting applications typically involve large databases with several millions of hours of video and audio, whereas traditional applications of image hashing typically focus on authenticating a smaller set of images. As a result, there are stringent requirements

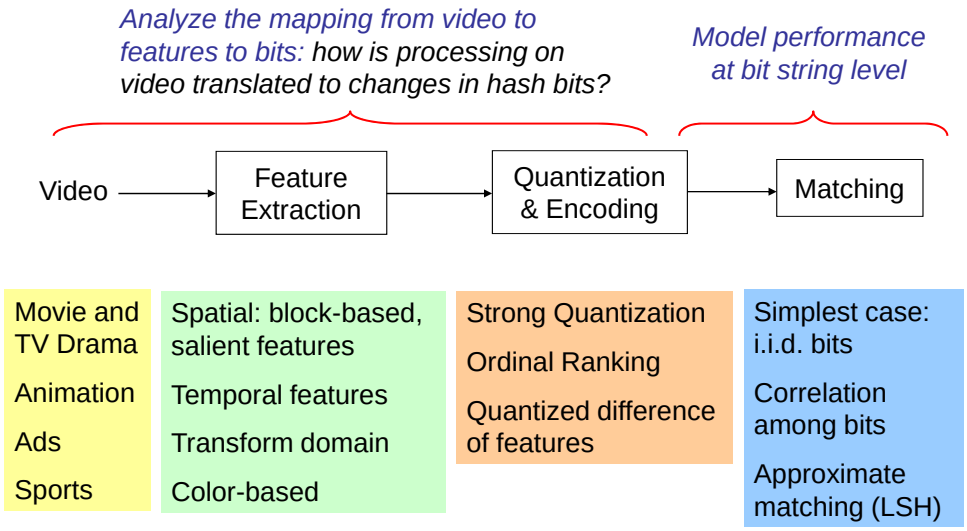

Figure 1. Framework for analysis of content fingerprinting.

on the computational complexity of fingerprinting schemes. Thus, theoretical studies focused on understanding content fingerprinting schemes are needed. In our previous work, we have analyzed the performance of binary fingerprints [4, 5], and examined optimal strategies for the designer and attacker from a game theoretic perspective [6]. We have also performed some preliminary analysis of the ordinal ranking module popularly used in fingerprinting [7]. In this paper, we propose a framework for systematic analysis of content fingerprinting schemes and summarize the analysis of some key modules under this framework.

2. PROPOSED FRAMEWORK FOR ANALYSIS OF CONTENT FINGERPRINTING

To facilitate the analysis of fingerprinting schemes, we propose the modular framework shown in Fig. 1. In a typical fingerprinting scheme, a set of robust and discriminative features are extracted from a given video and suitably quantized and encoded to obtain a compact representation of the features that is referred to as a fingerprint. Given a database of videos, this process is repeated for each video in the database to obtain a corresponding database of fingerprints. When a certain query video has to be identified, the fingerprint is computed and compared with the database of known fingerprints, represented by the matching block in Fig. 1.

Popular choices for the features extracted from the fingerprints include various spatial, temporal, transform domain or color-based features [1]. Such features are converted into a fingerprint by suitable quantization such as 1-bit quantization by comparing to a threshold, or through ordinal ranking [8]. Another popular option is to quantize the differences between successive feature components [9]. For comparing a fingerprint to a database, while exhaustive matching would be optimal in terms of performance, the computational requirements are often prohibitive. Instead, approximate search techniques such as Locality Sensitive Hashing (LSH) [10] are often employed in practice.

The overall robustness, accuracy of identification and the computational complexity is influenced by the particular choices for each of these blocks. For example, block-based spatial features such as the difference in average luminance can be computed efficiently, but are less robust to geometric distortions. On the other hand, the use of transform domain features such as the Fourier-Mellin Transform can provide robustness against geometric distortions [11] at a slightly higher computational cost. It has also been observed that the genre of the video under consideration, such as animations, sports videos, etc., can significantly impact the effectiveness of certain features.

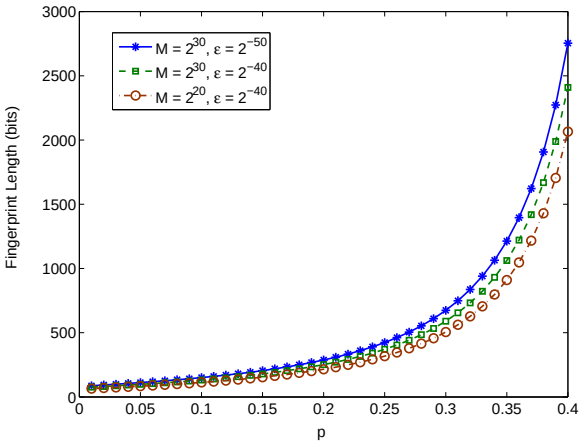

Figure 2. Fingerprint length N required to achieve a target probability of false alarm ϵ as a function of the noise probability p for database of size M.

To better understand the performance of various fingerprinting techniques, it is necessary to understand the impact of each of these individual modules on the overall system. Specifically, we are interested in understanding how various processing on different types of videos affects the features extracted, and how such changes in the features translate into changes in the fingerprints. These modifications in the fingerprints would in turn affect the matching accuracy of the identification system. In this work, we illustrate such a modular approach by focusing on some key modules of typical fingerprinting systems. Section 3 examines the performance bounds of binary fingerprints and Section 4 focuses on analyzing the impact of ordinal ranking on the robustness. Section 5 concludes the paper.

3. PERFORMANCE MODELING OF BINARY FINGERPRINTS

Binary fingerprints are often employed for content identification purposes as they can be stored and compared efficiently. We first focus our analysis on understanding the performance bounds of binary fingerprints and deriving guidelines for choosing system parameters to achieve desired performance. We briefly summarize the results for independent and identically distributed (i.i.d.) fingerprints followed by an analysis of the performance of correlated binary fingerprints.

3.1 Analysis of i.i.d. Equiprobable Binary Fingerprints

From a designer's perspective, it is beneficial to have binary fingerprints with components that are independent and equally likely to be 0 or 1, as each fingerprint bit then conveys the maximum amount of information. If the fingerprint bits are not independent with 0 and 1 equiprobable, they could be compressed to obtain a shorter string with i.i.d. bits, which is advantageous from a compactness and storage perspective. If the fingerprint bits are independent, an attacker cannot change a significant fraction of the bits at once by making small changes to the multimedia. Further, from a game-theoretic perspective, the optimal strategy for the designer is to choose fingerprint bits that are equally likely [6]. Hence, we first focus on analyzing the performance of i.i.d. bits with 0 and 1 equiprobable. With appropriate preprocessing and synchronization, distortions of the underlying video can be represented as additive noise in the fingerprint domain [2], which we model as flipping each fingerprint bit independently with a probability p.

Given a multimedia database of size M with corresponding fingerprints $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_M$, and a query multimedia document to be identified, the corresponding fingerprint $\mathbf{y}$ is obtained and compared with each fingerprint in the database. We model this as an $(M+1)$-ary hypothesis testing problem, with the null hypothesis representing the case that the query fingerprint $\mathbf{y}$ does not correspond to any fingerprint in the database, and the i^{th}

hypothesis H_i corresponding to the case that $\mathbf{y}$ is a noisy version of $\mathbf{x}_i$ [4]. Under this setting, it can be shown that to achieve a false alarm probability less than ϵ, a high probability of detection, and robustness against noise of strength p using fingerprints of length N, the fingerprint length must satisfy

$$\frac{1}{N}\log\frac{M}{\epsilon} < 1 - h(p), \tag{1}$$

where $h(p) = -p\log p - (1-p)\log(1-p)$ is the entropy function [4]. Figure 2 shows the relation between the fingerprint length and the desired robustness for given probability of false alarm and size of the database.

3.2 Analysis of Correlated Binary Fingerprints

In the previous subsection, we discussed the performance of i.i.d binary fingerprints from a detection-theoretic perspective, in terms of the probability of false positives and false negatives. As practical schemes often result in fingerprint bits with pairwise correlations, it is important to model and analyze their performance. In this section, we describe a Markov Random Field based model for fingerprints with correlated bits.

We illustrate our model using the example of a block-based video fingerprinting scheme that computes fingerprints from individual frames by extracting one bit from each block of the frame. For example, the fingerprint bit could be obtained by quantizing the average luminance within each block to 1-bit accuracy, or by quantizing the difference in the average luminance of adjacent blocks. We model the fingerprint bits using a Markov Random Field (MRF) as shown in Fig. 3(a). Each node in the MRF denotes one fingerprint bit and the edges between nodes capture local dependencies among the fingerprint bits. The energy function and the joint distribution for this MRF model are defined as

$$\begin{aligned} E_0(\mathbf{x}) &= -h\sum_i x_i - \eta\sum_{\langle j,k\rangle} x_j x_k, \\ p_0(\mathbf{x}) &= \frac{1}{Z_0}\exp(-E_0(\mathbf{x})) \end{aligned}$$

where the x_is denote the nodes of the MRF which may be ± 1 corresponding to the fingerprint bit being 0 or 1, respectively, $\langle j,k\rangle$ denotes those nodes that are joined by an edge and Z is a normalization constant. The parameters h and η control the marginal distribution and the correlation among the bits, respectively. The typical correlation observed among fingerprint bits is shown in Fig. 3(c) obtained by setting $h = 0$ and $\eta = 0.2$.

The noisy fingerprint bits are modeled using the MRF shown in Fig. 3(b), where the open circles represent the fingerprint bits of the undistorted frame and the filled circles represent the noise. The solid lines capture the local dependencies among the underlying fingerprint bits while the dashed and dotted lines capture the correlations among the noise bits. The dotted lines can capture dependencies between the fingerprint and noise bits. For simplicity, we will only consider the case where the noise bits are mutually correlated but independent of the fingerprint bits, implying that the parameters associated with the dotted edges in Fig 3(b) are zero. The energy function and the joint distribution of the noise is defined as

$$\begin{aligned} E_1(\mathbf{n}) &= -\alpha\sum_i n_i - \gamma\sum_{\langle j,k\rangle} n_j n_k, \\ p_1(\mathbf{n}) &= \frac{1}{Z_1}\exp(-E_1(\mathbf{n})) \end{aligned}$$

where n_is represent the noise bits and the parameters α and γ control the marginal distribution and the correlation of the noise bits, respectively.

The problem of matching a given fingerprint with a database can be considered as a multiple hypothesis testing problem [4], but may be well approximated by a series of binary hypothesis tests. Hence, as a first step, we consider a binary hypothesis test where given a particular query fingerprint and a database fingerprint, the

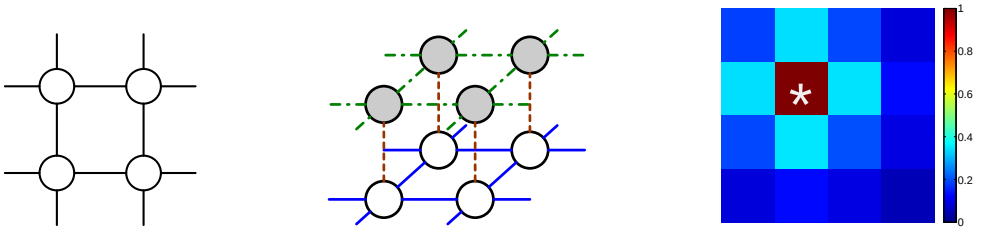

(a) (b) (c)

Figure 3. (a) MRF model for correlated binary fingerprints, (b) model for generation of noisy fingerprints and (c) the typical correlation observed between bits for a 4×4 sized model. The white asterisk indicates the bit under consideration.

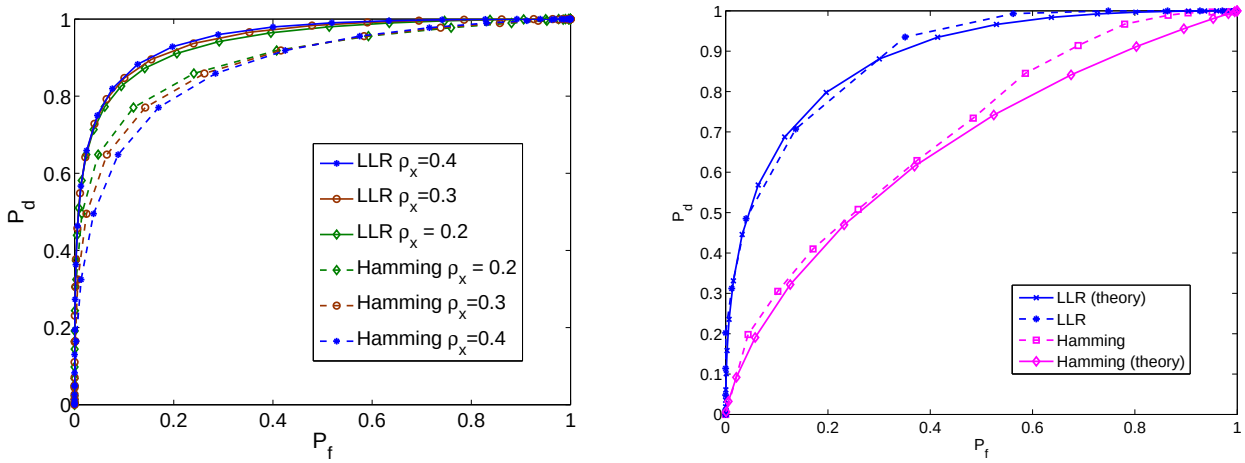

(a) (b)

Figure 4. (a) Influence of fingerprint correlation on detection performance using LLR and Hamming distance detectors. (b) Theoretical predictions shown by the solid lines agree well with the experimental results shown by the dashed lines.

detector aims to decide whether they are independent or not. Inspired by statistical physics, we use a density of state estimation technique [12] to estimate the probability of correct detection and false alarm. Details regarding this procedure may be found in [5].

Fig. 4(a) shows the receiver operating characteristic curves for the optimal Log-Likelihood Ratio (LLR) and Hamming distance detectors for different correlation among the fingerprint bits. The size of the fingerprint is 4×4, the probability of a noise bit being -1 (corresponding to a bit 1) is set to be 0.2 and the correlation among the noise bits is also 0.2. From the figure, we see that the optimal LLR detector is more robust to variations in the correlation among the bits as it compensates for this correlation. On the other hand, the performance of the suboptimal Hamming distance detector is affected by the correlation among the bits. Fig. 4(b) compares the theoretical predictions with experimental results obtained over a database of 1000 images with the noisy images generated by histogram equalization. The fingerprint was generated by dividing each image into 4×4 blocks and quantizing the average luminance of each block to 1-bit accuracy. From the figure, we see that the theoretical results agree with the experimental results.

The above analysis allows us to predict and quantify how changes in fingerprint bits affect the identification performance. It is also important to understand how changes in the feature values affect the fingerprint bits. The mapping from feature values to fingerprint bits is done via quantization and encoding. Ordinal ranking is one of the quantization modules that is commonly used for robust feature representation in content fingerprinting, although little analytic study has been made on its performance. In the next section, we provide theoretical analysis of ordinal ranking and examine its impact on the detection accuracy of fingerprinting schemes.

4. QUANTIFYING THE PERFORMANCE IMPACT OF ORDINAL RANKING

Ordinal ranking has been proposed as a technique to robustly represent a set of features by sorting them and recording only their ranks [8]. For example, the feature vector $[7.2, 4, 8, 6.3]$ is converted into a rank vector $[3, 1, 4, 2]$. As minor changes in the feature values would not impact the relative order, feature representation using ordinal ranking is believed to improve the robustness to noise, which has been confirmed by experiments. At the same time, as the ranking is global, strong local distortions that significantly change a single feature may alter all the ranks and cause significant changes in the overall fingerprint. In this section, we analyze the sensitivity of ordinal ranking to such local changes and quantify the impact of using ordinal ranking on the overall detection performance.

4.1 Sensitivity of Ordinal Ranking

We characterize the sensitivity of ordinal ranking in terms of the expected change in the rank of a given feature when the feature values are perturbed. Consider a block-based feature extraction scheme with N blocks. From each block B_i, $i = 1, 2, \ldots, N$, a real-valued scalar feature X_i is extracted and is normalized to have zero mean and unit variance. The rank of block B_i is denoted by Y_i, and $\{Y_i\}$ form a set taking values from 1 to N.

First, consider the simple case where the feature value from a single block is perturbed. Without loss of generality, we assume that a variational value $\Delta > 0$ is added to X_1, and the block ranks are recalculated and denoted by Z_i, $i = 1, 2, \ldots, N$. We define the rank difference of B_i with respect to variation Δ as $D_i(\Delta) = Z_i - Y_i$, and examine the expected value of rank difference, $\mathbb{E}[D_i(\Delta)]$. In practice, each feature value is usually an average or some combination of sample values from a block, and by the Central Limit Theorem, may be well approximated by a Gaussian as the number of values in a block is sufficiently large in practice. Assuming that each feature value is Gaussian-distributed with the correlation coefficient between X_i and X_j being $\rho_{i,j}$, it can be shown that the expected change in the rank of block B_i, $i \geq 2$, is given by

$$\mathbb{E}[D_i(\Delta)] = -\left(\Phi\left(\frac{\Delta}{\sqrt{2(1-\rho_{1,i})}}\right) - \frac{1}{2}\right), \tag{2}$$

and the expected difference in the rank of the first block B_1 is correspondingly

$$\mathbb{E}[D_1(\Delta)] = \sum_{i=2}^{N} \Phi\left(\frac{\Delta}{\sqrt{2(1-\rho_{1,i})}}\right) - \frac{N-1}{2}, \tag{3}$$

where $\Phi(\cdot)$ is the cumulative distribution function of the standard Normal distribution. As $\Phi(\cdot)$ is a monotonically increasing function, higher Δ and $\rho_{1,i}$'s imply larger sensitivity. An important observation is that the expected difference in B_1's rank can be seen as the sum of $(N-1)$ individual terms involving B_1 and one other B_i only.

The variance of the rank difference can also be derived similarly, but the expressions have been omitted due to space constraints. Further, these results can also be generalized to the case where multiple feature values are perturbed [7]. Fig. 5 shows the mean and the standard deviation of the change in the rank of a block as a function of the local variation Δ. From the figure, we observe that both the mean and the standard deviation increase as the correlation between the features increases.

4.2 Impact on the Detection Performance

We next compare the detection accuracy with and without ordinal ranking to understand the impact of ordinal ranking on the overall detection performance. Again, we consider the binary hypothesis test where given a query fingerprint and a fingerprint from the database, the detector aims to identify whether the two fingerprints correspond to different versions of the same video or two distinct unrelated videos. Without ordinal ranking, two feature vectors are matched by comparing their Euclidean distance to a threshold T. With ordinal ranking, the

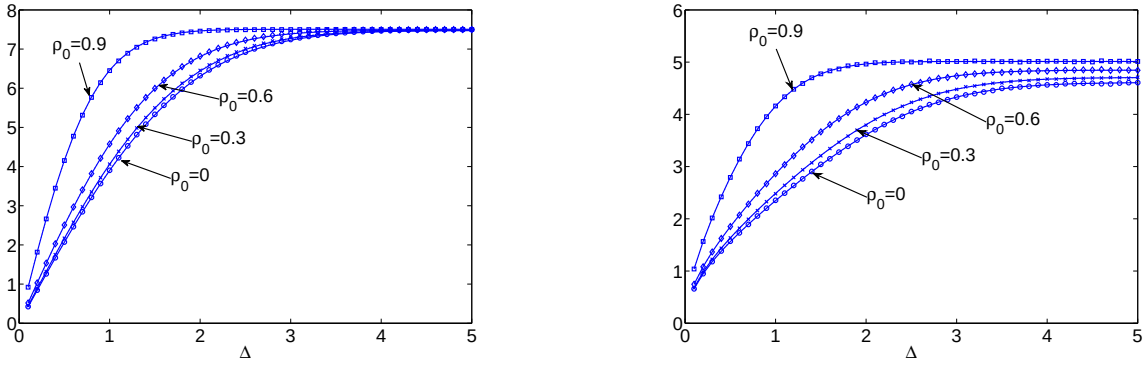

(a) $\mathbb{E}\{D_1(\Delta)\}$ (b) $\mathrm{Std}(D_1(\Delta))$

Figure 5. (a) Expected value and (b) Standard deviation of the change in the rank of a block as a function of the local variation Δ for various correlations.

feature vectors are converted into their ordinal rank vectors and the L_1 distance between the two rank vectors is compared to a threshold T'.

Suppose $\mathbf{p}$ and $\mathbf{q}$ are two independent fingerprint vectors generated according to the fingerprint distribution and $\tilde{\mathbf{p}}$ is a noisy version of $\mathbf{p}$. Without ordinal ranking, the probability of detection and false alarm are given by:

$$\begin{aligned} P_F &\triangleq \Pr\{\|\mathbf{p}-\mathbf{q}\|_2 \le T\}, \\ P_D &\triangleq \Pr\{\|\mathbf{p}-\tilde{\mathbf{p}}\|_2 \le T\}. \end{aligned} \tag{4}$$

With ordinal ranking, let $\mathrm{OR}(\mathbf{x})$ denote the ordinal rank vector of feature vector $\mathbf{x}$. The probability of detection and false alarm are then given by:

$$\begin{aligned} P_F^{\mathrm{OR}} &\triangleq \Pr\{\|\mathrm{OR}(\mathbf{p})-\mathrm{OR}(\mathbf{q})\|_1 \le T'\}, \\ P_D^{\mathrm{OR}} &\triangleq \Pr\{\|\mathrm{OR}(\mathbf{p})-\mathrm{OR}(\tilde{\mathbf{p}})\|_1 \le T'\}. \end{aligned} \tag{5}$$

For our analysis, we assume that each element of the feature vector is i.i.d Gaussian and is normalized to have zero mean and unit variance. The noise is modeled as i.i.d. zero mean Gaussian entries with a variance P. Under a Neyman-Pearson setting, we set the probability of false alarm to be ϵ and compare the probability of detection achievable with and without ordinal ranking. A higher probability of detection under the same probability of false alarm indicates higher robustness to noise. Here, we focus on the case when the fingerprint length is moderately large ($N \ge 10$ from our experiments).

Without ordinal ranking, the square of the detection statistic corresponds to the squared Euclidean distance between two Gaussian vectors with i.i.d. entries which is distributed according to a χ^2-distribution with N degrees of freedom. The probability of detection when the probability of false alarm is fixed to be ϵ is then given by

$$P_D = F_N\left(\frac{2F_N^{-1}(\epsilon)}{P}\right), \tag{6}$$

where $F_N(\cdot)$ denotes the cumulative distribution function (CDF) of a χ^2 distributed random variable with N degrees of freedom.

4.2.1 Detection Accuracy with Ordinal Ranking

For fingerprinting with ordinal ranking, when the two fingerprints are independent, the detection statistic is $\|\mathrm{OR}(\mathbf{p})-\mathrm{OR}(\mathbf{q})\|_1 = \sum_{i=1}^{N} |Y_i(\mathbf{p}) - Y_i(\mathbf{q})|$, where $Y_i(\mathbf{x})$ denotes the rank of the i^{th} element of $\mathbf{x}$. As $\mathbf{p}$ and $\mathbf{q}$

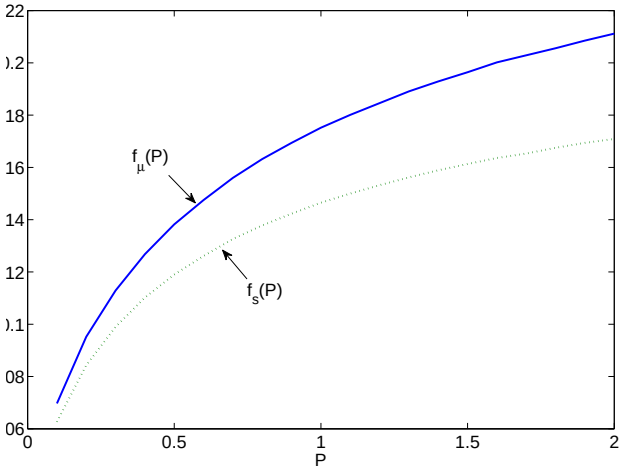

Figure 6. Numerical evaluations of $f_\mu(P)$ and $f_s(P)$.

are independent, OR($\mathbf{p}$) and OR($\mathbf{q}$) are two independent permutations of the vector $[1 \;\; 2 \;\; \cdots \;\; N]$. Denote the i^{th} entry of these two permutation vectors by r_i and t_i, respectively. Since the distribution of $\sum_{i=1}^{N} |r_i - t_i|$ is the same as that of $\sum_{i=1}^{N} |r_i - i|$, one of the permutation vectors can be fixed to be the identity. When N is sufficiently large, for $i \neq j$, $|r_i - i|$ and $|r_j - j|$ are only weakly correlated, and by the Central limit Theorem, the sum of many such terms, $\sum_{i=1}^{N} |r_i - i|$, can be well approximated by a Gaussian random variable. The mean and variance of this Gaussian distribution can be shown to be approximately $(N^2 - 1)/3$ and $N^3/20$, respectively.

Now we consider the distribution of $\|\text{OR}(\mathbf{p}) - \text{OR}(\tilde{\mathbf{p}})\|_1$. Following a similar argument, we can approximate the detection statistic by a Gaussian random variable $\mathcal{N}(\mu, s^2)$. The exact expressions for μ and s are difficult to evaluate, but from our experiments, we observe that both μ and s are linearly proportional to N: $\mu = N f_\mu(P)$ and $s = N f_s(P)$. Intuitively, both μ and s should increase as N increases, although the exact linear relations remain to be proved. Fig. 6 shows numerical estimates of f_μ and f_s as a function of the noise power P.

With the obtained distributions of the test statistics under the two hypotheses, we can now approximate the probability of detection for a given false alarm rate ϵ as:

$$P_D^{\text{OR}} \approx \Phi\left(\frac{2\sqrt{5N}[1 - 3f_\mu(P)] + 3\Phi^{-1}(\epsilon)}{6\sqrt{5} f_s(P)}\right), \tag{7}$$

where $\Phi(\cdot)$ is the cumulative distribution function of a Gaussian random variable. As $\Phi(\cdot)$ is an increasing function, when $1 - 3f_\mu(P) > 0$, or equivalently, $P < f_\mu^{-1}(1/3)$, the detection probability increases as N increases, implying that using more features gives higher robustness. On the other hand, if $P > f_\mu^{-1}(1/3)$, then the detection probability decreases as N increases for fingerprinting with ordinal ranking. However, as shown in Fig. 6, the value of the noise variance P at which $f_\mu(P) > 1/3$ is quite large and such large noise may not be a concern in practical applications. Eqn. (7) can be used as a guideline for determining the length of features required to achieve a desired P_D and P_F.

4.2.2 Comparison of P_D and P_D^{OR}

Fig. 7 compares the probability of detection with and without ordinal ranking, P_D^{OR} and P_D, for $N = 16$ and $N = 32$ as a function of noise variance P when the probability of false alarm $P_F = 10^{-3}$. The theoretical predictions for the probability of detection are obtained using Eqns. (6) and (7). From the figures, we see that for the probability of detection without ordinal ranking P_D, the theoretical and experimental results agree very

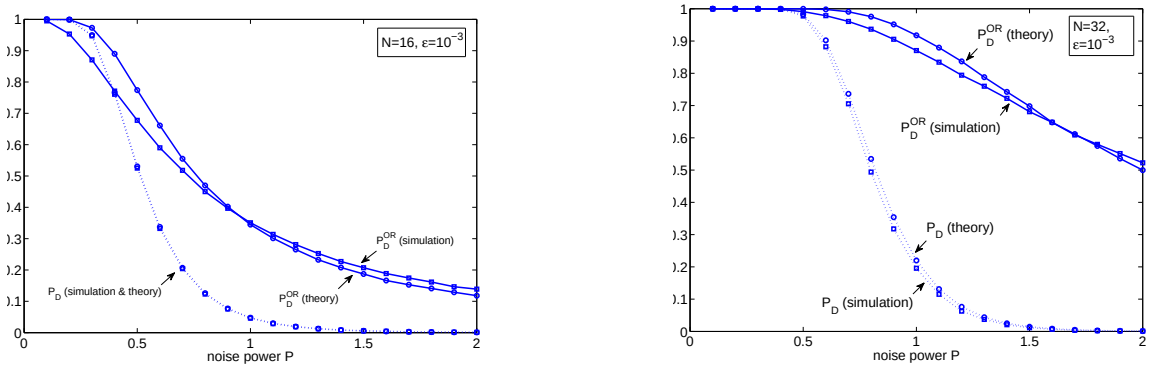

(a) N = 16 (b) N = 32

Figure 7. Comparison of P_D and P_D^{OR} for different values of P when (a) $N = 16$ and (b) $N = 32$.

well. For the probability of detection with ordinal ranking P_D^{OR}, the theoretical and experimental results are close, but there is a small discrepancy due to the approximations involved.

For $N = 32$, both the theoretical and simulation results indicate that ordinal ranking increases the robustness of the fingerprinting system and improves the probability of detection. For $N = 16$, we observe that at low noise power $P < 0.6$ the simulation results indicate that performance with ordinal ranking is actually lower than without ordinal ranking, which is not very well captured by the theoretical results. For practical applications where N is expected to be relatively large for each frame, such an effect will be small, and the theoretical expression would predict the trend of the detection accuracy reasonably well.

5. CONCLUSIONS

In this work, we have described a framework for modeling and analysis of content fingerprinting schemes using a modular approach and analyzed some key modules of fingerprinting systems. We first examined the performance of fingerprint matching at the bit string level and presented a bound on the performance achievable using fingerprints with i.i.d. equiprobable bits. We then described a Markov Random Field model for binary correlated fingerprints and examined the impact of the fingerprint correlation on the detection performance. Experimental results on an image database corroborated our theoretical predictions.

We also examined the impact of ordinal ranking on the performance of a fingerprinting system. We first presented results for quantifying the sensitivity of ordinal ranked features to local distortions. Our analysis shows that features with higher correlations are more susceptible to local variations when represented using their ordinal ranks. We then derived expressions for the probability of detection when ordinal ranking is used for encoding features. Comparing the performance attained using ordinal ranking with the performance obtained by using the raw features directly, we conclude that when the number of features is sufficiently large, using ordinal ranking indeed improves the robustness against global distortions.

REFERENCES

[1] J. Lu, "Video Fingerprinting for Copy Identification: from Research to Industry Applications," in *SPIE Media Forensics and Security*, Jan. 2009.

[2] E. McCarthy, F. Balado, G. Silvestre, and N. Hurley, "A Framework for Soft Hashing and its Application to Robust Image Hashing," in *IEEE Int. Conf. on Image Proc.*, Oct. 2004, pp. 397–400.

[3] S. Voloshynovskiy, O. Koval, F. Beekhof, and T. Pun, "Robust Perceptual Hashing as Classification Problem: Decision-theoretic and Practical Considerations," in *IEEE Workshop on Multimedia Signal Processing*, Oct. 2007, pp. 345–348.

[4] A. L. Varna, A. Swaminathan, and M. Wu, "A Decision-Theoretic Framework for Analyzing Binary Hash-based Content Identification Systems," in *ACM Workshop on Digital Rights Management*, Oct. 2008, pp. 67–76.

[5] A. L. Varna and M. Wu, "Modeling Content Fingerprints using Markov Random Fields," in *IEEE Int. Workshop on Information Forensics and Security*, Dec. 2009.

[6] ——, "Theoretical Modeling and Analysis of Content Identification," *IEEE Int. Conf. on Multimedia and Expo*, Jul. 2009.

[7] W.-H. Chuang, A. L. Varna, and M. Wu, "Modeling and Analysis of Ordinal Ranking in Content Fingerprinting," in *IEEE Int. Workshop on Information Forensics and Security*, Dec. 2009.

[8] R. Mohan, "Video Sequence Matching," in *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, vol. 6, Apr. 1998, pp. 3697–3700.

[9] J. Haitsma, T. Kalker, and J. Oostveen, "Robust Audio Hashing for Content Identification," in *Int. Workshop on Content-Based Multimedia Indexing*, Brescia, Italy, Sept. 2001.

[10] A. Gionis, P. Indyk, and R. Motwani, "Similarity Search in High Dimensions via Hashing," in *Int. Conf. on Very Large Databases*, 1999, pp. 518–529.

[11] A. Swaminathan, Y. Mao, and M. Wu, "Robust and Secure Image Hashing," *IEEE Trans. on Information Forensics and Security*, vol. 1, no. 2, pp. 215–230, June 2006.

[12] F. Wang and D. P. Landau, "Efficient, Multiple-Range Random Walk Algorithm to Calculate the Density of States," *Physical Review Letters*, vol. 86, no. 10, pp. 2050–2053, Mar. 2001.