Derivation of ROCs for Composite Fingerprints and Sequential Trimming

January 7, 2010

In this report, we derive the error probabilities for the fingerprint database search for sequential trimming and for composite fingerprints as described in [1]. All fingerprints have n pixels and there are N fingerprints in the database. Both n and N may be large (e.g., larger than 10^6). Formally, matching a query fingerprint to the fingerprint database is a multiple-hypothesis testing problem.

We assume that all database fingerprints are of the same quality (SNR), including the query fingerprint. Each estimated database fingerprint is modeled as $\mathbf{F} = \{\mathbf{F}[j]\}_j = \{x_j + \xi_j\}_j$, $j = 1, \ldots, n$, where $\xi_j \sim N(0, \sigma^2)$ is WGN and x_j is the true fingerprint, which is a fixed realization of a Gaussian $N(0, 1-\sigma^2)$. In other words, the fingerprints are normalized to have sample variance of 1. Also realize that each fingerprint estimate is a realization of a standard Gaussian $N(0, 1)$. Furthermore, we also assume that fingerprints of two different cameras are independent.

The direct brute-force search computes the normalized correlation, ρ, between each database fingerprint $\mathbf{F}$ and the query fingerprint $\mathbf{X}$:

$$\rho(\mathbf{F}, \mathbf{X}) = \frac{\sum_{j=1}^{n}(\mathbf{F}[j] - \overline{\mathbf{F}})(\mathbf{X}[j] - \overline{\mathbf{X}})}{\sqrt{\sum_{j=1}^{n}(\mathbf{F}[j] - \overline{\mathbf{F}})^2}\sqrt{\sum_{j=1}^{n}(\mathbf{X}[j] - \overline{\mathbf{X}})^2}},$$

where the bar denotes the sample mean.

1 Sequential trimming

Here, we derive the distribution of the normalized correlation ρ for non-matching and matching fingerprints. From the independence of fingerprints assumption, using the Central Limit Theorem (CLT) $\rho(\mathbf{F}_i, \mathbf{F}_{i'}) \sim N(0, 1/n)$ for $i \neq i'$. For fingerprints trimmed to length k, $\rho(\mathbf{F}_i, \mathbf{F}_{i'}) \sim N(0, 1/k)$.

To derive the distribution under H_1, we assume $\{x_j + \xi_i\}$, $\{x_j + \eta_j\}$ are two noisy versions of the same fingerprint and of the same quality (i.e., $Var[\xi] = Var[\eta] = \sigma^2$). Note that

$$SNR = \frac{\sum x_j^2}{\sum \xi_j^2} = \frac{\sum x_j^2}{\sum \eta_j^2} = \frac{1-\sigma^2}{\sigma^2},$$

which implies

$$(1 + 1/SNR)^{-1} = 1 - \sigma^2. \tag{1}$$

The correlation is (all sums are over k elements of the trimmed fingerprints):

$$\begin{aligned}
\rho(x_j + \xi_i, x_j + \eta_i) &\approx \frac{\sum x_j^2 + \sum x_j(\xi_j + \eta_j) + \sum \xi_j \eta_j}{\sum x_j^2 (1 + 1/SNR)} \\
&\approx (1 + 1/SNR)^{-1} + \frac{1}{k(1 + 1/SNR)(1 - \sigma^2)} \left(\sum x_j(\xi_j + \eta_j) + \sum \xi_j \eta_j \right) \\
&= 1 - \sigma^2 + \frac{1}{k} \left(\sum x_j(\xi_j + \eta_j) + \sum \xi_j \eta_j \right).
\end{aligned}$$

Above, we used (1) and $\frac{1}{k} \sum x_j^2 \approx 1 - \sigma^2$.

Because the variance of the product of two independent Gaussians is the product of their variances, by the CLT:

$$\frac{1}{\sqrt{k}}\left(\sum x_j(\xi_j+\eta_j)+\sum \xi_j\eta_j\right) \xrightarrow{d} N\left(0,(1-\sigma^2)\times 2\sigma^2+\sigma^4\right)=N(0,2\sigma^2-\sigma^4),$$

with means that

$$\rho \sim N(1-\sigma^2,\sigma_c^2),$$
$$\sigma_c^2=\frac{2\sigma^2-\sigma^4}{k}.$$

Choosing the threshold t for the detector, the search proceeds by computing the correlation for each database fingerprint. We now have several options how to make a decision from all N computed correlations.

[Conservative detector] The detection is successful only when all $N-1$ non-matching correlations are $\rho < t$ and $\rho \geq t$ for the matching fingerprint. The probability of false alarm and the probability of detection are thus:

$$\begin{aligned} P_{\mathrm{FA}} &= 1-(1-Q(t\sqrt{k}))^N, \\ P_{\mathrm{D}} &= (1-Q(t\sqrt{k}))^{N-1}Q((t-1+\sigma^2)/\sigma_c), \end{aligned} \tag{2}$$

where $Q(x)$ is the probability that a standard normal random variable exceeds x. The ROC for this detector has a curious shape. With decreasing threshold t, $P_{\mathrm{FA}} \to 1$, as expected, but the detection $P_{\mathrm{D}} \to 0$ instead of to 1 because the false alarms will prevent the detection from being successful. In practice, the search may inspect all candidate fingerprints whose correlation was above the threshold and then double check the match on full-length fingerprints. This will increase the search time, however. In order to compare the trimming and composite fingerprints fairly, we use this conservative detector.

[Maximal correlation] Alternatively, the detected fingerprint may be chosen as the one with the highest correlation. A complete analysis of this case is quite complex due to the involved order statistics. Reference [2] (Section 4.5.3) has the derivation of the minimal total probability of error $P_{\mathrm{e}}=(P_{\mathrm{FA}}+P_{\mathrm{D}})/2$ for the Bayesian case under equal error costs and equal priors. This detector will have a higher P_{D} for a given P_{FA} than the conservative detector.

2 Composite fingerprints

The composite fingerprint is ($s=(N+1)/2$):

$$\mathbf{M}=\left\{\frac{1}{\sqrt{s}}\left(\sum_{l=1}^{s} x_j^{(l)}+\xi_j^{(l)}\right)\right\}=\left\{\frac{1}{\sqrt{s}}x_j+\xi'_j\right\},$$

where we singled out one specific fingerprint $x_j=x_j^{(1)}$ for our next analysis. The factor $1/\sqrt{s}$ is to make $\mathbf{M}$ of unit variance. The variance $Var[\xi']=\frac{s+\sigma^2-1}{s}$ because

$$\begin{aligned} \sum x_j^2+\sum \xi_j^2 &= \frac{1}{s}\sum x_j^2+\sum \xi_j'^2 \\ 1-\sigma^2+\sigma^2 &= \frac{1-\sigma^2}{s}+Var[\xi']. \end{aligned} \tag{3}$$

The $SNR'=\frac{\frac{1}{s}\sum x_j^2}{\sum \xi_j'^2}$ for the composite fingerprint is obtained by dividing (3) by $\sum x_j^2$:

$$\begin{aligned} 1+\frac{1}{SNR} &= \frac{1}{s}+\frac{1}{s}\frac{1}{SNR'} \\ 1+1/SNR' &= s(1+1/SNR). \end{aligned}$$

Under H_0,

$$\rho \sim N(0,1/n)$$

because we are again correlating two independent Gaussians.

To derive the distribution of ρ under H_1, whenever a query fingerprint $\mathbf{F} = \{x_j + \xi_j\}_j$ is in the composite fingerprint $\mathbf{M} = \{\frac{1}{\sqrt{s}}x_j + \xi'_j\}_j$:

$$\rho(\mathbf{F}, \mathbf{M}) = \rho\left(x_j + \xi_j, \frac{1}{\sqrt{s}}x_j + \xi'_j\right) \approx \frac{\frac{1}{s}\sum x_j^2 + \sum x_j(\xi'_j + \frac{1}{\sqrt{s}}\xi_j) + \sum \xi'_j\xi_j}{\frac{1}{s}\sum x_j^2\sqrt{1+1/SNR}\sqrt{1+1/SNR'}} = \frac{1-\sigma^2}{\sqrt{s}} + \frac{1}{\sqrt{n}}\Xi,$$

where the noise term is

$$\Xi = \frac{\sqrt{s}}{\sqrt{n}}\left(\sum x_j(\xi'_j + \frac{1}{\sqrt{s}}\xi_j) + \sum \xi'_j\xi_j\right).$$

Above, all sums are over $j = 1, \ldots, n$ and thus $\sum x_j^2 \approx n(1-\sigma^2)$. Again, by the CLT, the noise term is approximated by a Gaussian

$$\begin{aligned} \Xi &\xrightarrow{d} N\left(0, (1-\sigma^2)(s+\sigma^2-1+\sigma^2) + \sigma^2(s+\sigma^2-1)\right) \\ &= N\left(0, s+\sigma^2-1+\sigma^2(1-\sigma^2)\right) \\ &= N(0, 2\sigma^2-\sigma^4+s-1). \end{aligned}$$

Thus, the correlation when a fingerprint is present in a composite fingerprint is

$$\begin{aligned} \rho &\sim N\left(\frac{1-\sigma^2}{\sqrt{s}}, \sigma_s^2\right), \\ \sigma_s^2 &= \frac{2\sigma^2-\sigma^4+s-1}{n}. \end{aligned}$$

For a fixed threshold t, the composite detector gives the correct answer only when $\rho \geq t$ for all composite fingerprints containing the query fingerprint. We note that $\rho < t$ if $\mathbf{X}$ is not in the database. Thus, the error rates are:

$$\begin{aligned} P_{\mathrm{FA}} &= 1-(1-Q(t\sqrt{n}))^m, \\ P_{\mathrm{D}} &= \frac{1}{N}\sum_{d=1}^{m}\binom{m}{d}P_{\mathrm{D}}^{(0)}(m-d)P_{\mathrm{D}}^{(1)}(d), \end{aligned}$$

where, $m = \log_2 N$ and

$$\begin{aligned} P_{\mathrm{D}}^{(0)}(m-d) &= \left(1-Q(t\sqrt{n})\right)^{m-d}, \\ P_{\mathrm{D}}^{(1)}(d) &= Q\left(\frac{t-\frac{1-\sigma^2}{\sqrt{s}}}{\sigma_s}\right)^d. \end{aligned} \tag{4}$$

Note that for large s and n, $\sigma_s^2 \approx \frac{s-1}{n}$. The ROC curve also bends towards zero for $P_{\mathrm{FA}} \to 1$ for the same reason as in the previous section on trimming.

3 Experiments

We now compare the performance of sequential trimming and composite fingerprints by plotting the corresponding ROCs. For a fair comparison, the length of the trimmed fingerprint, k, has been always adjusted so that both search methods lead to the same expected search time. The composite approach needs to compute $\log_2 N$ full-length correlations instead of N for the direct brute-force search. The search is thus $N/\log_2 N$ times faster. To obtain the same speed for trimming, the value of k was set to $k = n/(N/\log_2 N)$. In all experiments below, $n = 1,920,000$, which is the native resolution of iPhone camera sensors.

Table 1 shows the typical and critical values of the maximum probability of detection $\max_{P_{\mathrm{FA}}} P_{\mathrm{D}}(P_{\mathrm{FA}})$ for a range of N, and a few selected values of ρ for each method. First, notice that the search based on trimming always

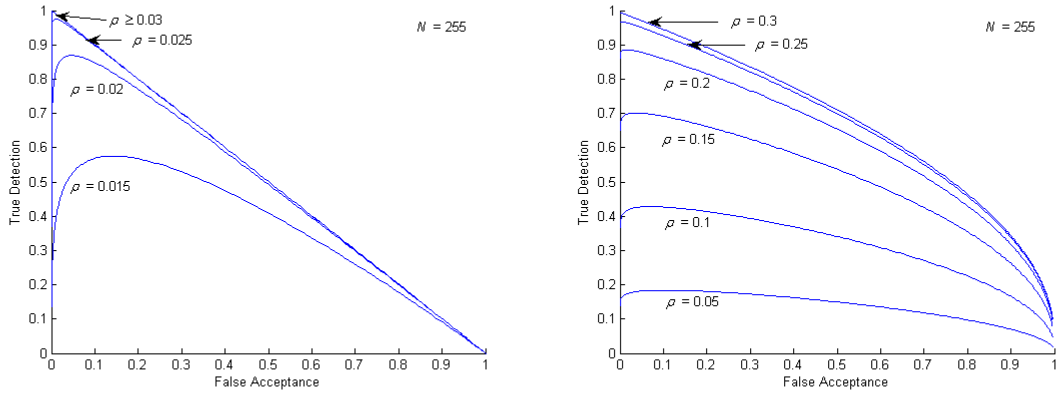

Figure 1: The ROC curve $P_{\mathrm{D}}(P_{\mathrm{FA}})$ for $N = 255$, for various values of the fingerprint correlation ρ. Left: trimming, Right: composite.

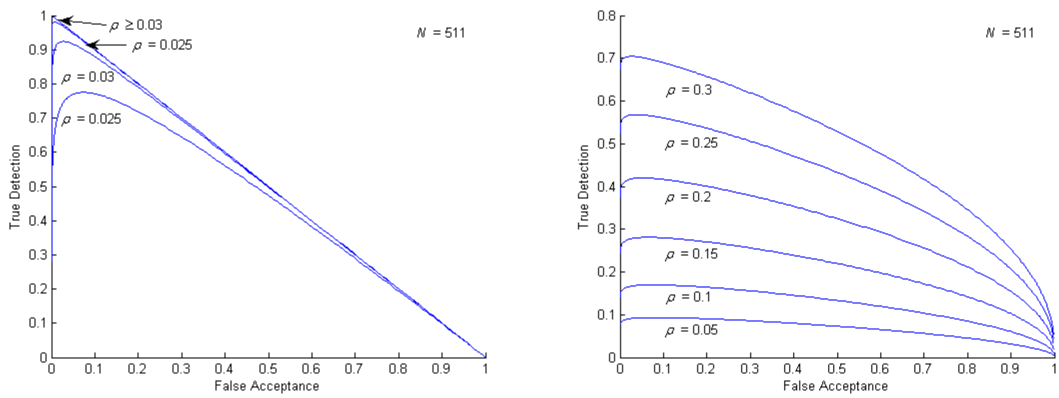

Figure 2: The ROC curve $P_{\mathrm{D}}(P_{\mathrm{FA}})$ for $N = 511$, for various values of the fingerprint correlation ρ. Left: trimming, Right: composite.

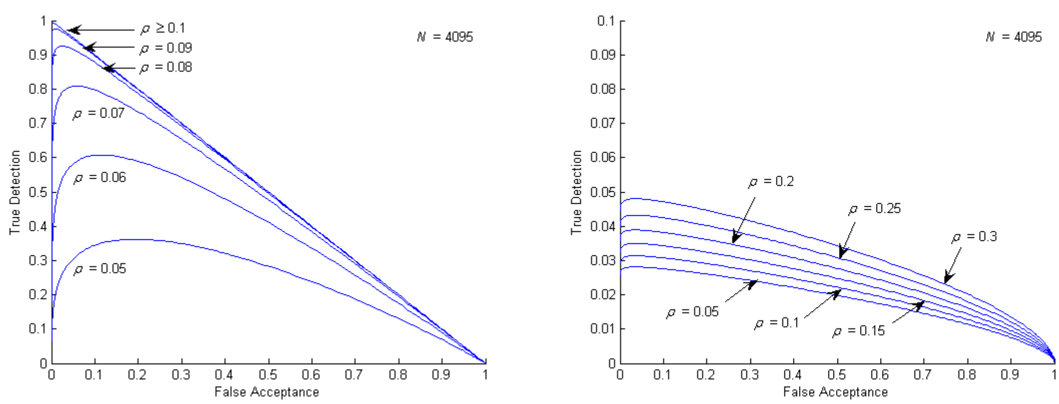

Figure 3: The ROC curve $P_{\mathrm{D}}(P_{\mathrm{FA}})$ for $N = 4095$, for various values of the fingerprint correlation ρ. Left: trimming, Right: composite.

			max P_{D}	
N	k	ρ	Trimming	Composite
127	106494	0.01	0.51	0.11
		0.02	0.99	0.16
		0.15	1	0.99
255	60429	0.02	0.87	0.09
		0.03	1	0.12
		0.30	1	0.99
511	60429	0.02	0.51	0.06
		0.04	1	0.08
		0.60	1	0.99
1023	18784	0.03	0.59	0.05
		0.05	0.99	0.06
		0.60	1	0.70
2047	10321	0.04	0.52	0.05
		0.06	0.95	0.06
		0.07	0.99	0.07
4095	5627	0.06	0.60	0.03
		0.10	0.99	0.03
		0.99	1	0.17
10,000	10,000	0.05	0.70	0.02[1]
		0.07	0.98	0.02[1]
100,000	10,000	0.05	0.54	N/A
		0.08	0.99	N/A
100,000	50,000	0.02	0.36	N/A
		0.03	0.93	N/A
		0.04	1	N/A

Table 1: Maximum achievable probability of detection P_{D} for various combinations of the database size N, trimmed fingerprint length k, and correlation between fingerprints, $\rho = 1 - \sigma^2$.

has a higher P_{D} than the search that uses composite fingerprints. While composites fail in the sense that the $\max P_{\mathrm{D}}$ falls off quickly for $N > 511$, trimming delivers $\max P_{\mathrm{D}} \approx 1$ for values of the correlation that are important for fingerprint search in practice even for a relatively small k ($k/n \ll 1$), independently of the database size N. Note that the maximum P_{D} is typically achieved for a rather small P_{FA}. Moreover, with increasing $\max P_{\mathrm{D}}$ the P_{FA} decreases towards zero (see Figures 1–3).

References

[1] M. Goljan, J. Fridrich, and T. Filler. Managing a large database of camera fingerprints. In E. J. Delp, J. Dittmann, and A. M. Alattar, editors, *Proceedings SPIE, Electronic Imaging, Media Security and Forensics XII*, volume 7541, San Jose, CA, January 18–20, 2010.

[2] S. M. Kay. *Fundamentals of Statistical Signal Processing, Detection Theory*, volume II. Prentice Hall, 1998.

[1] For the number of pixels $n = 7,525,500$, which matches the complexity of trimming at $k = 10,000$.