

Towards Quantifying the Entropy of Fingervein Patterns across Different Feature Extractors

Vedrana Krivokuća
Idiap Research Institute
Rue Marconi 19, Martigny, Switzerland
vedrana.krivokuca@idiap.ch

Sébastien Marcel
Idiap Research Institute
Rue Marconi 19, Martigny, Switzerland
sebastien.marcel@idiap.ch

Abstract

This paper makes a first attempt at quantifying the entropy of fingervein patterns that have been extracted using three different state-of-the-art feature extractors, on two publicly-available fingervein databases. We show that the resulting entropy is dependent upon both the feature extractor and database, implying that a universal estimate of fingervein entropy would be misleading. We also discuss how our entropy results can be applied towards more meaningful evaluations of the security and privacy of fingervein template protection schemes. Our open-source implementation of the entropy estimation on a publicly-available fingervein recognition system will help the research community to both validate our findings and build upon our work.

1. Introduction

Fingervein recognition is increasing in popularity, with companies such as M2SYS, Morpho-Safran, Hitachi, and NEC having already deployed fingervein recognition technologies in practice. Unfortunately, the development of strategies to secure our fingervein features during storage and transmission in fingervein recognition systems is seriously lagging. A recent review paper on biometric template protection [1] showed that, in the last decade, only 1% of all biometric template protection research accounts for fingerveins. Our paper increases that 1% by contributing towards fingervein template protection research.

Two of the most important requirements of a biometric template protection scheme are that it should preserve both the security of the underlying recognition system and the privacy of the system's users. Preservation of security

means that, after the biometric templates are protected, they should still contain enough information to remain useful for distinguishing between different identities. Preservation of privacy means that the amount of information about the unprotected biometric template that is leaked (revealed) by the protected template is insufficient to recover the original template or a close approximation of it. Although various methods have been adopted in the literature to evaluate the security and privacy of biometric template protection schemes, we believe that such an analysis cannot be complete without taking into account the amount of information contained in the original, unprotected biometric template. In other words, if we do not know how much information we started with, how can we provide a complete, meaningful analysis on the amount of information lost or gained as a result of applying the biometric template protection scheme? In this paper, we present the results of a preliminary investigation into the amount of information contained in fingervein patterns, or the *entropy* of fingervein patterns, when those patterns have been extracted using three state-of-the-art feature extractors.

To the best of our knowledge, estimates of the entropy of fingervein patterns do not exist in the literature. A few papers (e.g., [2, 3]) analyse the individuality of fingerveins in terms of their ability to distinguish between different identities. While this proves that fingerveins can be used for recognition, estimates of the information content of fingervein patterns are not provided. Moreover, the analysis is generally based on proprietary databases and, as far as we know, the underlying implementation is not freely available. This makes it difficult to reproduce the results or implement a fingervein template protection scheme on top of the baseline system. We fill this gap by estimating the entropy of fingervein patterns on two public databases, using three state-of-the-art feature extractors, for which the full implementation is publicly available. This way, the research community can implement fingervein template protection schemes on top of our adopted baseline system, and our entropy results as well as our open-source entropy estimation implementa-

tion can then be applied in analysing and comparing those template protection schemes in a more meaningful way.

The remainder of this paper is structured as follows. Section 2 outlines the methodology used to estimate fingervein entropy. Section 3 presents and analyses our results. Section 4 discusses how our results can be used to evaluate fingervein template protection schemes, and we also mention a perceived drawback of the adopted entropy estimation method. Section 5 provides concluding remarks and ideas for future work in this direction.

2. Methodology

When estimating the entropy of fingervein features, what we are essentially trying to do is to approximate the amount of discriminatory information contained in fingervein patterns, where “discriminatory information” refers to features used to differentiate between different fingers. This means that the estimated entropy must be heavily influenced by the underlying fingervein features. Since different feature extraction methods are likely to produce differences in the extracted fingervein patterns, it is reasonable to assume that there are also likely to be differences in the perceived information content of the extracted fingervein patterns depending on the adopted feature extractor. For example, an extractor that extracts clean fingervein patterns would likely produce patterns with a different entropy to those produced by an extractor that extracts noisy fingervein patterns.

Taking the above points into account, it becomes evident that a *universal* estimate of fingervein entropy would be rather misleading, since we cannot guarantee that the entropy of fingervein features extracted in one way would be the same as the entropy of fingervein features that have been extracted in another way. For this reason, we believe that estimates of fingervein entropy should be system-specific, and we prove this point by showing the differences between our entropy estimates for three different automated feature extractors and two publicly-available fingervein databases.

Section 2.1 discusses the adopted fingervein databases, Section 2.2 outlines the fingervein feature extraction process, and Section 2.3 explains how we estimated the entropy of fingervein patterns from the extracted features.

2.1. Databases

We used two public fingervein databases for our investigation: VERA¹ and UTFVP². VERA consists of two images for each of 110 people’s left and right index fingers, which makes up 440 images in total. UTFVP consists of four images for each of 60 people’s left and right index, ring and middle fingers, which makes up 1,440 images in total.

¹<https://www.idiap.ch/dataset/vera-fingervein>

²<http://scs.ewi.utwente.nl/downloads/show,FingerVein/>

Both databases were captured using the same imaging device, but with slightly different acquisition conditions. For more information on these two databases, refer to [4, 5].

2.2. Feature Extraction

To extract the fingervein patterns from the images in the databases, we used the open-source bob.bio.vein package³ implemented using Idiap’s Bob toolbox [6, 7]. First, the finger in each image is located and horizontally aligned as per [8, 9]. Next, the vein pattern is extracted from the finger images using three state-of-the-art extractors: Wide Line Detector (WLD) [9], Repeated Line Tracking (RLT) [10], and Maximum Curvature (MC) [11]. The output of each extractor is a binary image, in which white pixels represent the fingervein pattern and black pixels represent the background (see Figure 1).




Figure 1: Fingervein patterns extracted using three different feature extractors on image 001.L.1 from VERA and image 0001.1.1 from UTFVP fingervein database.

We concatenated the rows in each binary image to generate fingervein feature vectors. Table 1 presents the sizes of the binary images and feature vectors.

Database	Extractor	Image Size	Feature Vector Size
VERA	WLD	62×162	10,044
	RLT	156×405	63,180
	MC	260×675	175,500
UTFVP	WLD	94×164	15,416
	RLT	234×409	95,706
	MC	390×682	265,980

Table 1: Sizes of the binary fingervein pattern images (pixels) and corresponding fingervein feature vectors (bits) generated by each extractor, for VERA and UTFVP databases.

Due to the specific characteristics of each extractor as well as the initial sizes of the images in the two databases, all the feature vector dimensionalities in Table 1 are different. It is differences like these that support our focus in this paper, which is to show that entropy estimates depend on

³<https://pypi.python.org/pypi/bob.bio.vein>

system parameters such as the adopted acquisition device (and thus the resulting database) and feature extractor.

2.3. Entropy Calculation

The *entropy* of a fingervein feature vector essentially refers to the amount of *discriminatory information* contained in that feature vector (i.e., the information that helps us distinguish between two fingers). There are a number of different approaches that could be used to estimate the fingervein entropy. For the investigation presented in this paper, our entropy estimate consisted of two steps. Sections 2.3.1 and 2.3.2, respectively, explain these two steps.

2.3.1 Step 1: Number of Independent Bits

Since our fingervein feature vectors are binary, the first step in our entropy estimation was inspired by the method Daugman used in [12] to estimate the number of independent bits in binary IrisCodes. The idea is to compute the Hamming distance (HD) between every pair of unrelated feature vectors to obtain a distribution of the HDs. The HD between two binary vectors is the proportion of disagreeing corresponding bits. The resulting distribution is then fitted to a binomial distribution with a certain number of degrees of freedom, which is computed using the HD distribution's mean and standard deviation. The degrees of freedom estimates the number of independent bits in each binary feature vector, which in turn provides an approximation of the discriminative capabilities of the underlying features.

So, we began by computing the HD between every pair of *different* (i.e., from different fingers) fingervein feature vectors, for each of our three extractors and each of the two databases. There was a total of 192,720 HDs for the VERA database and 2,067,840 for UTFVP. We then plotted the distributions of the resulting HDs, and calculated the mean, standard deviation, and degrees of freedom for each distribution. The number of degrees of freedom was computed using Equation 1 [12], where p represents the mean of the HD distribution and σ denotes its standard deviation:

$$N = \frac{p(1-p)}{\sigma^2} \quad (1)$$

An important characteristic of any binomial distribution is that the underlying experiment consists of a set of *independent* Bernoulli trials, where the outcome of each trial can be either 1 (“success”) or 0 (“failure”). In our case, one “trial” is a comparison between two bits of two fingervein feature vectors (so there would be B trials for every pair of B -bit feature vectors), and a “success” is a mis-matched pair of bits between the two feature vectors (so the proportion of successes is the HD). The number of degrees of freedom of a binomial distribution tells us the number of *independent trials* used to compute the distribution. So, if we can approximate our HD distribution by an N -degrees-of-freedom

binomial distribution, we can estimate that our HD distribution consists of N independent trials, which implies that there are approximately N independent bits in the comparison of two fingervein feature vectors. Consequently, we may assume that there are N bits of discriminatory information between our binary fingervein feature vectors. This can be used to estimate the entropy of our fingervein patterns, as discussed in Section 2.3.2.

2.3.2 Step 2: Shannon Entropy

A binomial distribution can be thought of as the distribution resulting from repeatedly tossing a coin N times and tallying up the proportion of times that the coin lands on “Heads” for each set of N trials. If the coin is *fair*, then we can expect it to land on “Heads” approximately 50% of the time, in which case the mean of the distribution (or the average probability of “success”) would be $p = 0.5$. If the coin is *unfair*, then we can expect the mean of the distribution to be $p > 0.5$ (if the probability of getting “Heads” is greater than the probability of getting “Tails”) or $p < 0.5$ (if “Heads” is less likely than “Tails”).

Let us now apply this same analogy to our fingervein HD distribution. In Section 2.3.1, we outlined the rationale behind attempting to approximate our HD distribution by a binomial distribution with N degrees of freedom. If such an approximation were valid, this would tell us that our HD distribution (i.e., the amount of difference between the fingervein feature vectors of different fingers) is distributed equivalently to runs of N tosses of a coin [12]. If the mean of our HD distribution was around 0.5, then it would be equivalent to runs of N tosses of a *fair* coin, and if the mean was more or less than 0.5 then we could approximate the distribution by runs of N tosses of an *unfair* coin.

This analogy can be used to estimate the Shannon entropy of our fingervein patterns. The Shannon entropy of a single coin toss with a probability of success (“Heads”) p and a probability of failure (“Tails”) $q = 1 - p$ is calculated using Equation 2, and the total Shannon entropy, in bits, of N coin tosses is calculated using Equation 3:

$$H = -p \log_2 p - q \log_2 q \quad (2)$$

$$H_T = NH \quad (3)$$

If the coin is *fair*, then $p = q = 0.5$ and thus $H = 1$ and $H_T = N$. If the coin is *unfair*, then $p \neq q$ and thus $H < 1$ and $H_T < N$.

So the second step in our entropy estimation was to calculate the Shannon entropy as per Equations 2 and 3, replacing N with the number of degrees of freedom (effective number of “coin tosses”), p with the mean of the HD distribution, and q with $1 - p$. The result of Equation 2 is an estimate of the average amount of information contained in

a single bit, and the outcome of Equation 3 is an estimate of the average number of bits needed to represent a fingervein feature vector for recognition purposes, which in turn provides an estimate of the amount of information we expect to gain from a comparison of two different fingervein feature vectors (when that comparison is based on the HD). Following the same line of thought, we could then estimate the total number of *unique* fingervein feature vectors to be 2^{H_T} .

The procedure outlined in Sections 2.3.1 and 2.3.2 was followed to estimate the entropy of fingervein patterns for each of the three feature extractors (WLD, RLT, and MC) and both databases (VERA and UTFVP). Our implementation of the entropy estimation is available at the following URL: <https://pypi.python.org/pypi/bob.paper.isba2018-entropy>. The results are presented and analysed in Section 3.

3. Results

This section presents our results for the entropy of fingervein patterns, when the fingervein images come from two different databases (VERA and UTFVP) and the fingervein patterns are extracted using three different feature extractors (WLD, RLT, and MC). Figure 2 shows the resulting HD distributions overlaid with the corresponding binomial distributions, and the mean, standard deviation, degrees of freedom and entropy for each distribution are summarised in Table 2.

DB	Extractor	p	σ	N	H	H_T
VERA	WLD	0.31	0.017	723	0.89	647
	RLT	0.37	0.021	547	0.95	519
	MC	0.08	0.007	1,413	0.42	588
UTFVP	WLD	0.24	0.017	594	0.79	469
	RLT	0.27	0.020	518	0.84	437
	MC	0.06	0.005	1,958	0.33	647

Table 2: Mean (p), standard deviation (σ), degrees of freedom (N), entropy per independent bit (H) and total entropy of N independent bits (H_T) for the HD distributions.

A number of important observations may be drawn from Figure 2 and Table 2. Section 3.1 discusses how well the HD distributions fit their corresponding binomial distributions, Section 3.2 considers how the HD distributions compare across the three feature extractors and two databases, Section 3.3 observes trends in the number of degrees of freedom, Section 3.4 discusses the entropy estimates, and Section 3.5 compares our estimates of the fingervein entropy to an estimate of the entropy of IrisCodes from [12].

3.1. Binomial Distribution Approximations

It is clear from the plots in Figure 2 that all of our HD distributions are well-approximated by their corresponding




Figure 2: HD distributions (light grey), with the corresponding binomial distributions overlaid (black), for the WLD, RLT, and MC feature extractors on the VERA and UTFVP fingervein databases. The HD distributions were normalised to better visualise the binomial distribution fit. The histogram bin width is 0.001.

N -degrees-of-freedom binomial distributions (i.e., the binomial distributions computed using the mean and standard deviation of the HD distributions). This suggests that the amount of difference between our binary fingervein feature vectors is indeed equivalent to runs of N tosses of a coin.

For example, the HD distribution for fingervein features extracted using the WLD extractor on the VERA database has a mean of 0.31 and 723 degrees of freedom. We may think of this as being equivalent to the probability distribution of obtaining “Heads” when tossing an unfair coin 723 times, when the probability of that coin landing on “Heads” is approximately 0.31. In other words, we can compare the underlying HD distribution to a binomial distribution with 723 degrees of freedom and a mean of 0.31. Similar conclusions can be drawn for the other extractors.

We may deduce, therefore, that our binary fingervein feature vectors consist of N independent bits, which can be used to recognise one finger from another in terms of the HD between the corresponding fingervein patterns.

3.2. HD Distributions for Different Extractors and Databases

From Figure 2, we can see that the HD distributions differ depending on which feature extractor is used. For example, while the WLD and RLT extractor distributions appear to have roughly the same shape and to be located at approximately the same position (for each database separately), the MC distribution is located much further to the left and has a much narrower shape. These differences may be due to the fact that the MC fingervein patterns are thinner and perhaps cleaner than the WLD and RLT patterns (e.g., see Figure 1). Consequently, there are likely to be more correlated background pixels (thereby producing a lower HD mean) and the calculated HDs are likely to be more consistent across multiple fingervein comparisons (thereby producing a small standard deviation). Considering these differences, it seems reasonable to conclude that fingervein entropy estimates should be extractor-specific.

Figure 2 also shows differences between the HD distributions for the same extractor across the VERA and UTFVP databases. In particular, while the shape of the distributions across the two databases is more or less the same (for each extractor separately), the UTFVP distributions appear to be left-shifted versions of their VERA counterparts. Consulting Table 2, we can see that the means (p) of the UTFVP extractors are indeed smaller than the p s of the corresponding extractors for the VERA database. One reason for this is probably due to the different image sizes for the same extractor across the two databases (see Table 1). For example, from Figure 1 we can see that UTFVP’s fingervein pattern images seem to have a larger amount of background information (black pixels) compared to VERA’s fingervein pattern images. This means that when two UTFVP fingervein feature vectors are compared, there is likely to be a larger number of matching bits (i.e., the background pixels), which would push the p of the HD distribution to the left. Cropping the region of interest (ROI) and discarding the background area would probably provide a closer comparison between the two databases, and we intend to do this as part of our future work; however, in this paper we simply use the outputs of the feature extractors (i.e., the binary fingervein patterns as illustrated in Figure 1) as they are.

3.3. Trends in Degrees of Freedom

From Table 2, it is evident that for both databases MC fingervein features have the most degrees of freedom (i.e., largest N) and RLT features have the least. This suggests that MC fingervein patterns consist of the greatest number of independent bits that can be used to discriminate between different fingers, when that discrimination is based on the HD, and RLT features contain the least.

We can also see that the standard deviation (σ) has much more influence on the resulting N than does the mean (p).

For example, for both databases the MC distribution has the smallest p and RLT the largest, yet RLT fingervein patterns have a significantly smaller N than MC features. This is because the σ of the MC distribution is much smaller than that of the RLT distribution. Considering Equation 1, we can see that this trend is expected due to the fact that σ is squared, so any change in σ will have a greater effect than p on N .

Finally, all three extractors have significantly fewer degrees of freedom (i.e., much smaller N) than the total number of bits in the fingervein feature vectors (see Table 1). Since N gives us an estimate of the number of *independent* bits in a binary vector, we may conclude that there is a considerable amount of correlation between the pixels in a binary fingervein pattern image and thus between the bits in the corresponding feature vector. If *all* the bits in the feature vectors were independent, we would expect N to be equal to the dimensionality of the corresponding feature vector.

3.4. Entropy Estimates

The means (p) of all our HD distributions are less than 0.5, implying that the amount of difference between fingervein patterns from different fingers can be approximated by runs of N tosses of an *unfair* coin. If p was 0.5, this would imply that each of the N independent bits is equally likely to be 1 or 0, in which case we could conclude that each of the N independent bits contains 1 full bit of information. We would thus expect the total amount of information contained in our feature vectors, or the entropy of our feature vectors, to be equal to N . This, however, is not the case for our experiments, as we can see from the fact that all the H estimates in Table 2 are below 1, suggesting that each independent bit contains less than 1 bit of information. Consequently, the total information entropy of our fingervein features, H_T , is observed to be lower than the corresponding N estimates, as expected. The most striking difference between the N and H_T estimates can be observed for the MC extractor. This is probably due to the extremely small p , which suggests that although MC fingervein patterns do have a large number of independent bits, those bits are perhaps more predictable than the bits in WLD or RLT fingervein features, meaning that each MC pattern bit carries less information than a WLD or RLT pattern bit (as can be seen from the H estimates).

Considering the H estimates from Table 2, we can see that, for both databases, RLT fingervein feature vectors contain the most information per independent bit, and MC features contain the least. As expected from Equation (2), this is directly related to the trend observed in the HD distribution means, i.e., the closer the p is to 0.5, the larger the H . This implies that RLT features are the closest to a uniform distribution in terms of the proportion of black (background) and white (vein) pixels. Conversely, MC features

contain the greatest proportion of background pixels, which would be expected to result in the smallest average HD between two MC feature vectors.

The trend in the estimated H values follows the exact opposite trend to that observed for the estimated N values. This is because N is much more influenced by σ than by p , whereas H depends on p alone. Consequently, the trend in our N values follows the σ trend. This implies that while RLT features might have the greatest average amount of information per bit (H), the MC features are much more stable in terms of being able to achieve a tighter HD distribution. The smaller the variation between the HDs, the more difficult it becomes to represent *different* feature vectors that produce an HD in that range, which is why we need more degrees of freedom or more independent bits (N) per MC feature vector compared to RLT or WLD feature vectors.

The estimated values for the total amount of entropy of our fingervein patterns, H_T , show mixed results. While RLT features were found to have the smallest H_T for both databases, the largest H_T was obtained for WLD features on the VERA database and MC features on the UTFVP database. If we considered the N values alone, we would conclude that MC features carry the most discriminative information, whereas the H estimates may lead us to conclude that the RLT extractor produces the most discriminative fingervein patterns. The H_T estimates, however, show that it is important to take both N and H into account when attempting to quantify the overall amount of discriminatory information available in a fingervein pattern. Moreover, it is clear that both the feature extractor and the database play a role in the entropy estimate, thereby suggesting that a universal entropy measure may be misleading.

We could use our results for H_T to estimate the total number of *unique* fingervein identities (for a particular feature extractor and database) as 2^{H_T} . Then, to make sure that all impostors in our fingervein recognition system are rejected, we should ensure that $FMR < \frac{1}{2^{H_T}}$.

3.5. Comparison to IrisCodes

Finally, it is worth noting that, if we convert the 249 degrees of freedom calculated for IrisCodes in [12] to an estimate of total entropy, H_T , using the same technique as that presented in this paper, we would obtain an entropy of 249 bits, which is lower than all the H_T estimates in Table 2. This may be partly due to the fact that these IrisCodes consist of only 2,048 bits, while our fingervein feature vectors have a much higher dimensionality (see Table 1). Furthermore, perhaps the IrisCodes underwent a larger amount of noise reduction prior to feature extraction, thereby reducing the amount of perceived information in the final feature vectors. However, in both cases, the entropy estimates were based on *features output by the adopted feature extractor*. So, since our features are those features that would be used

for recognition purposes, there is a justifiable amount of fairness in this comparison, which suggests that it is worthwhile investigating further whether fingervein features may indeed be better at distinguishing between different identities than the IrisCodes from [12]. Part of this work should consider how the intra-class variation of fingervein patterns compares to the intra-class variation of IrisCodes, and the effect of this on the resulting entropy estimation. Section 4.2 discusses the intra-class variation point in more detail.

4. Discussion

In this section, we discuss how our results for the entropy of fingervein patterns can be applied towards evaluating the security and privacy of fingervein template protection schemes, and we mention a perceived drawback of the adopted entropy estimation method. Sections 4.1 and 4.2 deal with these two points of discussion, respectively.

4.1. Security and Privacy Analysis

An analysis of the *security* of a fingervein template protection scheme should consider how well fingervein entropy is preserved in the protected templates. If the protected templates remain in binary format, we could apply the same analysis to estimate the entropy of our protected templates, then calculate the difference between the entropies of the protected and unprotected fingervein templates, where the unprotected template is simply the extracted binary fingervein pattern. If the entropy of the protected templates is x bits less (more) than the entropy of the unprotected templates, this would imply that we *lose (gain)* approximately x bits of discriminatory information as a result of applying the template protection scheme. Although the FMR also provides a measure of the security of a biometric recognition system, an analysis of the security in terms of the entropy would be helpful since we could remove the reliance on a particular matching threshold.

An analysis of the *privacy* of a biometric template protection scheme generally considers the amount of information that the protected template leaks about the original data. Such an analysis would be more meaningful if we knew how much information was contained in the original data. For example, if we know that our original fingervein feature vector (i.e., the set of extracted features) contains k bits of information and the protected template leaks j bits (where $j \leq k$), then we know that the remaining (unrevealed) amount of information is approximately $k - j$ bits. Say that we need a minimum of n bits of information to decide whether or not two fingervein templates come from the same finger. Then, as long as $k - j \leq n$, the amount of leaked information should be insufficient to reliably link the protected template to the original template. Consequently, we could conclude that the leaked information cannot reveal which original feature vector the protected template came

from. Our results on the entropy of fingervein patterns can be applied in exactly this way to evaluate the privacy of fingervein template protection schemes.

A potential point of contention in the reader’s mind at this stage may be the following. Consider a fingervein template protection scheme that simply quantises the extracted binary fingervein pattern, such that the result is a smoothed version of this pattern (e.g., a smoothed version of the images in Figure 1). In this case, the protected template is likely to have a *lower* entropy than the original template, which would indicate a *loss of information* as a result of applying the template protection scheme. However, what if the protected fingervein pattern is actually a cleaner, more ‘noise-free’ version of the original template? This would seem to indicate that the only ‘information’ that was lost is noise, and noise is irrelevant for recognition purposes. This might lead the reader to think that it is wrong to conclude that the lower entropy of the protected template is a bad thing, thereby questioning the use of entropy as a measure of the security and privacy of template protection schemes as outlined in the previous two paragraphs. However, recall that the point of this paper was to estimate the entropy of fingervein features extracted by a particular feature extractor. In this case, we must assume that the extracted features are indeed *features* of the fingervein pattern, and thus *information* as opposed to noise. Any noise that is present in the extracted features is an imperfection of the feature extractor and this is something we must put up with if we want to use that extractor for recognition purposes. Consequently, if we assume that the fingervein pattern extracted by the feature extractor is all “information”, then it makes sense to conclude that a loss in entropy as a result of template protection is indeed a bad thing since it results in a loss of “information” (not noise). We believe, therefore, that entropy is a useful measure of the security and privacy of biometric template protection schemes, as long as we are clear about how exactly we define “features” or “information”.

Finally, since our analysis in Section 3 showed that the entropy of fingervein patterns depends upon both the feature extractor and fingervein database, we cannot guarantee that our results provide a universal measure of fingervein entropy; indeed, one of our aims was to illustrate this. So, we recommend that the results be applied for analysing fingervein template protection schemes that are tested in the same (or similar) experimental environment (i.e., same database, same extractor). Since our code is open-source, we encourage other researchers to build upon our implementation and thus help ensure that different fingervein template protection schemes are more easily comparable.

4.2. Drawback of Entropy Estimation

While the adopted entropy estimation method is certainly correct from the data compression or source coding

point of view [13], unfortunately there is a perceived drawback. As mentioned in Section 2.3.1, the adopted entropy estimation considers *inter*-class (i.e., between feature vectors from *different* fingers) HDs only. Consequently, the resulting entropy estimate is based on the underlying assumption that the *intra*-class variation (i.e., the HD between feature vectors from multiple samples of the *same* finger) is zero [13], which effectively implies that two feature vectors are expected to match either fully or not at all [14]. Since there is always some variation between multiple samples from the same biometric in practice, we are unlikely to get zero intra-class variation. As a result, the entropy reported for fingerveins in this paper is likely to be an over-estimate of the actual entropy expected in practice. This is because our entropy estimate gives us an indication of the amount of “surprise” in witnessing an HD of h between two feature vectors. The more times that h is witnessed, the less the surprise of encountering it again. If the only distance encountered in our intra-class HD distribution is 0, then this would not interfere with any of the inter-class HD distances, since we would expect all those distances to be greater than 0. If, however, we have some non-zero intra-class distances, there is a greater chance of encountering the same distance in the inter-class HD distribution, which would lessen the surprise of witnessing that distance and would thus lower the estimated entropy of fingervein patterns.

In an attempt to extend the idea of using entropy as a measure of biometric information while more practically incorporating both inter- and intra-class variation, several authors have adopted the *mutual information* or *relative entropy* approach based on the Kullback-Leibler divergence [15, 14, 16]. We are currently working on supplementing the findings presented in this paper with estimates of the amount of information in fingervein patterns based on the relative entropy measure.

5. Conclusions and Future Work

This paper presented an investigation into the entropy of fingervein patterns for three state-of-the-art feature extractors (WLD, RLT, and MC) on two publicly-available fingervein databases (VERA and UTFVP). Our entropy analysis was based on estimating the number of independent bits in automatically-extracted binary fingervein feature vectors, then computing the Shannon entropy for the binomial distribution that represents those independent bits. We observed that the entropy varies depending on the feature extractor and database used, so we concluded that it is best to provide system-specific estimations of fingervein entropy. From our results, it appears that the entropy of fingervein patterns is higher than that for IrisCodes. So, it may be worthwhile to conduct an investigation into determining whether fingervein patterns are indeed more reliable than IrisCodes for distinguishing between different identities. The results

we obtained for the entropy of fingervein patterns can be directly applied towards more meaningful evaluations of the security and privacy of fingervein template protection schemes. Since our entropy estimation code is open-source and is implemented on top of an open-source fingervein recognition system, this will enable the research community to build upon our work, thereby helping to speed up the advancement of fingervein template protection research.

Although we believe that the investigation presented in this paper is a good start towards estimating the entropy of fingervein patterns (and hope that our readers consider the entropy of this paper to be high), a lot more remains to be done before we can be sure of the robustness of this work. We currently have three plans for future work in this direction. Firstly, we think it would be interesting to repeat the experiments in this paper after removing the excess background regions from the extracted fingervein patterns. We expect that this would increase the mean of the resulting HD distributions, so it would be useful to see how this affects the entropy estimates. Secondly, we would like to investigate the entropy across different finger types. In this paper, each finger was treated as an individual entity, since it has been proven that different fingers from the same person can be seen as separate identities, so we did not distinguish between different finger types. Nevertheless, it would be interesting to see whether there exist significant differences between the amount of information, or the entropy, contained in different finger types and to quantify these differences. Finally, due to the fact that the adopted entropy measure does not take into account intra-class variation, we plan to investigate alternative ways of estimating fingervein entropy (primarily the relative entropy approach) and compare the results to the estimates presented in this paper.

6. Acknowledgements

We would like to acknowledge the Swiss Center for Biometrics Research and Testing, as well as the SWAN project, for funding this research. We also extend our thanks to André Anjos and Olegs Nikisins from Idiap Research Institute for their work on the fingervein recognition system framework on which our entropy estimation was based.

References

- [1] M. Sandhya and M.V.N.K. Prasad. *Biometric Template Protection: A Systematic Literature Review of Approaches and Modalities*, pages 323–370. Springer International Publishing, Cham, 2017.
- [2] Y. Ye, H. Zheng, L. Ni, S. Liu, and W. Li. A study on the individuality of finger vein based on statistical analysis. In *2016 International Conference on Biometrics (ICB)*, pages 1–5, June 2016.
- [3] T. Yanagawa, S. Aoki, and T. Ohyama. Human finger vein images are diverse and its patterns are useful for personal identification. *MHF Prepr. Ser.*, 12:17, 2007.
- [4] B.T. Ton and R.N.J. Veldhuis. A high quality finger vascular pattern dataset collected using a custom designed capturing device. In *2013 International Conference on Biometrics (ICB)*, pages 1–5, June 2013.
- [5] M. Vanoni, P. Tome, L. El-Shafey, and S. Marcel. Cross-database evaluation using an open finger vein sensor. In *2014 IEEE Workshop on Biometric Measurements and Systems for Security and Medical Applications (BIOMS) Proceedings*, pages 30–35, Oct 2014.
- [6] A. Anjos, M. Günther, T. De Freitas Pereira, P. Korshunov, A. Mohammadi, and S. Marcel. Continuously reproducing toolchains in pattern recognition and machine learning experiments. In *International Conference on Machine Learning (ICML)*, August 2017.
- [7] A. Anjos, L. El Shafey, R. Wallace, M. Günther, C. McCool, and S. Marcel. Bob: a free signal processing and machine learning toolbox for researchers. In *20th ACM Conference on Multimedia Systems (ACMMM)*, Nara, Japan, October 2012.
- [8] E.C. Lee, H.C. Lee, and K.R. Park. Finger vein recognition using minutia-based alignment and local binary pattern-based feature extraction. *International Journal of Imaging Systems and Technology*, 19(3):179–186, 2009.
- [9] B. Huang, Y. Dai, R. Li, D. Tang, and W. Li. Finger-Vein Authentication Based on Wide Line Detector and Pattern Normalization. In *2010 20th International Conference on Pattern Recognition*, pages 1269–1272, Aug 2010.
- [10] N. Miura, A. Nagasaka, and T. Miyatake. Feature extraction of finger-vein patterns based on repeated line tracking and its application to personal identification. *Machine Vision and Applications*, 15(4):194–203, 2004.
- [11] N. Miura, A. Nagasaka, and T. Miyatake. Extraction of finger-vein patterns using maximum curvature points in image profiles. *IEICE TRANSACTIONS on Information and Systems*, 90(8):1185–1194, 2007.
- [12] J. Daugman. How iris recognition works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, Jan 2004.
- [13] Y. Sutcu, E. Tabassi, H.T. Sencar, and N. Memon. What is biometric information and how to measure it? In *2013 IEEE International Conference on Technologies for Homeland Security (HST)*, pages 67–72, Nov 2013.
- [14] K. Takahashi and T. Murakami. A measure of information gained through biometric systems. *Image and Vision Computing*, 32(12):1194–1203, 2014.
- [15] A. Adler, R. Youmaran, and S. Loyka. Towards a Measure of Biometric Information. In *2006 Canadian Conference on Electrical and Computer Engineering*, pages 210–213, May 2006.
- [16] Y. Sutcu, H.T. Sencar, and N. Memon. How to Measure Biometric Information? In *2010 20th International Conference on Pattern Recognition*, pages 1469–1472, Aug 2010.