Towards Practical Cancelable Biometrics for Finger Vein Recognition

Christof Kauba [1a], Emanuela Piciucco[b], Emanuele Maiorana[b,*], Marta Gomez-Barrero[c], Bernhard Prommegger[a], Patrizio Campisi[b] and Andreas Uhl[a]

[a]Department of Computer Sciences, University of Salzburg, Jakob-Haringer-Str. 2, 5020 Salzburg, Austria
[b]Department of Industrial, Electronics and Mechanical Engineering, Roma Tre University, Italy
[c]Fakultät Wirtschaft, Hochschule Ansbach, Germany

ARTICLE INFO

Keywords:
biometrics
template protection scheme
cancelable biometrics
finger vein recognition
Bloom filters

ABSTRACT

Biometrics has nowadays become a preferred solution for systems requiring secure authentication. However, the usage of biometric characteristics raises significant concerns regarding personal data protection and privacy. Several template protection schemes have been therefore proposed to conceal the employed identifiers, while still ensuring the ability to efficiently recognise users. In this paper, we propose and analyse three different approaches generating cancelable templates from finger vein features. A thorough analysis of the considered methods is conducted to investigate their impact on the achievable recognition performance, as well as their security in terms of renewability and unlinkability. Furthermore, a specific attack is designed to evaluate the irreversibility of the protection scheme providing the best recognition performance.

1. Introduction

The adoption of biometric recognition for security purposes is constantly increasing in various practical applications within the field of human-machine systems, including border control, logical and physical access control, ATM cash withdrawal, and many more. The main reasons for such rapid spread are the enhanced customer convenience and the improved security this technology offers with respect to traditional authentication methods, such as those based on passwords or tokens. In fact, while these latter can be forgotten or stolen, it is not possible to loose or forget a biometric characteristic. Additionally, biometric characteristics are much more difficult to be fraudulently copied or forged than standard identifiers.

Despite the aforementioned advantages, the use of biometric data in recognition systems may also involve severe security and privacy concerns. Due to their uniqueness, biometric characteristics can allow an attacker to track the activities of a subject whose characteristics have been registered in different domains [40]. Moreover, compromised

[1]Please cite this work as: Christof Kauba, Emanuela Piciucco, Emanuele Maiorana, Marta Gomez-Barrero, Bernhard Prommegger, Patrizio Campisi and Andreas Uhl, Towards Practical Cancelable Biometrics for Finger Vein Recognition, Elsevier Information Sciences, in press, 2021

*Corresponding author
ckauba@cs.sbg.ac.at (C.K.); emanuela.piciucco@uniroma3.it (E. Piciucco); emanuele.maiorana@uniroma3.it (E. Maiorana); marta.gomez-barrero@hs-ansbach.de (M. Gomez-Barrero); bprommeg@cs.sbg.ac.at (B. Prommegger); patrizio.campisi@uniroma3.it (P. Campisi); uhl@cs.sbg.ac.at (A. Uhl)
http://www.wavelab.at/member-ckauba.shtml (C.K.); http://www.wavelab.at/member-bprommeg.shtml (B. Prommegger)
ORCID(s): 0000-0002-2716-1360 (C.K.); 0000-0001-7236-3511 (E. Piciucco); 0000-0001-9448-1316 (E. Maiorana); 0000-0003-4581-5353 (M. Gomez-Barrero); 0000-0003-4002-2602 (B. Prommegger); 0000-0002-1923-2739 (P. Campisi); 0000-0002-5921-8755 (A. Uhl)

biometric data cannot be used anymore, thereby further limiting the already small number of usable biometric identifiers that a subject can use [7]. Since biometric information cannot be revoked and reissued as it happens for disclosed passwords or stolen keys, proper countermeasures should be taken in order to address the aforementioned issues.

Biometric template protection (BTP) schemes have been therefore proposed to ensure the secure and private handling of biometric data during the authentication process. In general, these methods modify the original biometric template with the aim of generating an alternative representation in a protected feature space, with no information leakage about the original sample. The comparison procedure is then carried out in this secure domain, thereby protecting the data during the whole recognition process.

According to the ISO/IEC 24745 standard [17], a properly defined BTP scheme should satisfy the following properties:

- *irreversibility*: given a protected template, it should not be possible to reconstruct the original biometric sample;
- *renewability*: from a given biometric sample, it should be possible to issue multiple protected templates;
- *unlinkability*: given two protected templates, generated form the same biometric information and stored in different systems, it should not be feasible to determine that they belong to the same subject;
- *performance*: using a BTP scheme should not significantly degrade the system recognition performance. Moreover, the recognition performance should not be sensitive to the parameters specifying the employed template protection step [38].

BTP schemes are categorized into two main classes: *biometric cryptosystems* and *cancelable biometrics* approaches. The former class can be further separated into *key-binding* methods, whose aim is to secure a cryptographic key by means of biometric data and vice versa [14], and *key-generating* approaches, which derive a cryptographic key from biometric data [42]. In contrast, cancelable biometric methods apply a key-dependent transformation function to the biometric data or templates to be secured. *Salting* approaches are defined using invertible transformations, with system security thus relying on the secret storage of the employed key. Conversely, *non-invertible transformations* can be applied either to the samples or to the original templates [18]. This latter category is of great importance, since it typically allows template comparison to be performed in the protected domain while using the same techniques employed in the unprotected scenario. Actually, techniques based on non-invertible transformations have been already proposed for several characteristics, such as fingerprint [44], face [3], iris [33], palmprint [25], and online signature [27], among others.

Cancelable biometrics is the focus of the present work, where the security of templates generated applying non-invertible transforms to samples of an emerging biometric modality, that is, the vascular patterns of human fingers, is evaluated.

Finger vein biometrics is receiving an ever increasing interest from both industry and the research community because of its convenience in the acquisition procedure, its robustness to presentation attacks, and its high recognition performance [29]. The imaging of subcutaneous vein patterns is feasible thanks to the haemoglobin capability of absorbing near infra-red (NIR) light: a camera sensitive to the 800-900 nm range produces images where blood vessels appear darker than the remaining body parts when illuminated with NIR radiation.

Cancelable biometric schemes for vein patterns have been first proposed in [15], where a Fourier-like transform over a finite field has been used on finger vein images, with template comparison performed through correlation-based distance metrics. Similarly, a hashing/binary filtering approach, based on the application of an alignment robust scheme in combination with index-of-maximum hashing, has been used for finger vein template protection in [21]. In addition, methods relying on the fusion with other characteristics have been also proposed, e.g., fusing finger vein patterns with fingerprint minutiae to design a cancelable multi-biometric system [45].

In this context, we present a thorough benchmark of several cancelable biometrics techniques applied to finger vein patterns. In more detail, building upon the authors' previous work in [32], we evaluate the effectiveness of three distinct approaches, namely *block remapping*, *image warping*, and *Bloom filters* in the feature domain. The two former methods can be applied as well as in the feature domain and to any image-based biometric characteristic (i.e., not restricted to binary templates), as in the case of face [36] and iris [12]. Their use for finger veins has been also proposed in our previous work [22], by applying them to vein patterns in the image domain.

One of the main advantages of Bloom-filter-based template protection, with respect to other approaches, is its applicability to different biometric modalities represented through binary templates, including iris [37], face [9], or fingerprint [1]. Being able to apply a single method to protect different image-based templates leads to a second advantage, that is, the feasibility of implementing a multi-biometric feature level fusion [10], possibly resorting to user-friendly combinations such as face and iris, or fingerprint and finger vein. In particular, Gomez-Barrero et al. presented in [10] a general method to extract Bloom-filter templates from any binary, fixed-length template, so to use them jointly. In the experiments carried out for the different biometric characteristics in the aforementioned works, it can also be observed that, when considering templates protected using Bloom filters, there is typically only a minimal biometric recognition performance loss with respect to the usage of unprotected templates. Actually, as for the vast majority of BTP schemes, small alignment issues have to be handled during pre-processing to manage the most challenging samples, as it may happen for facial images [10]. In this regard, iris recognition represents a convenient scenario for the application of Bloom filters, since misalignments of the original samples can be handled by processing normalized iriscodes column-wise, and then discarding the information about which column originated a given bit in the protected template, without the need for computing multiple shifted versions of protected templates [37]. Finally, it should be noted that the Bloom filter extraction and comparison steps are fast, and that the generated

templates are sparse, that is, the BTP method also preserves computational efficiency.

In summary, the main contributions of the present article are:

- thorough evaluation and benchmark, in terms of recognition performance, irreversibility, unlinkability, and renewability, of three distinct cancelable biometrics approaches applied to finger vein patterns;
- application of the block-remapping and image-warping BTP schemes to feature representations of vein patterns. Differently from [32, 6], we consider six different feature representations of finger vein samples in order to evaluate the most appropriate one to be used in the protected biometric recognition system;
- use of Bloom filters to protect finger vein patterns. In contrast to [10], we apply Bloom filters directly to binary vein images, instead of using vein minutiae-based templates;
- proposal of a pre-alignment method for improving the recognition performance attainable by the employed finger vein cancelable biometrics;
- exploitation of a specific attack, based on a square jigsaw puzzle solver algorithm, to quantify the irreversibility of the block remapping approach.

2. Finger Vein Recognition

The standard finger vein recognition processing pipeline includes: acquisition of the input vein image, image pre-processing, feature extraction, and template comparison.

2.1. Pre-processing

In order to localise the area containing the relevant finger vein patterns, the region of interest (ROI) is first detected within the acquired image [24]. Possible misalignments, due to different finger positioning in distinct acquisitions, are subsequently compensated with a normalisation step [16]. It should be noted that only a coarse alignment can be typically achieved, thereby needing to further correct errors by performing shifts in both directions during the comparison step. More importantly, such issues are much more severe when dealing with templates transformed through block-based BTP approaches like the ones we consider here, thereby raising the need for a pre-alignment stage, as discussed in Section 3.4.

Finally, image quality is further improved by enhancing the contrast of the ROI image, and non-uniform illumination is compensated by applying Contrast Limited Adaptive Histogram Equalization (CLAHE) [50], High Frequency Emphasis Filtering (HFEF) [49], and Circular Gabor Filtering (CGF) [48]. An example of a finger vein sample and its corresponding pre-processed image can be found in Figure 1.

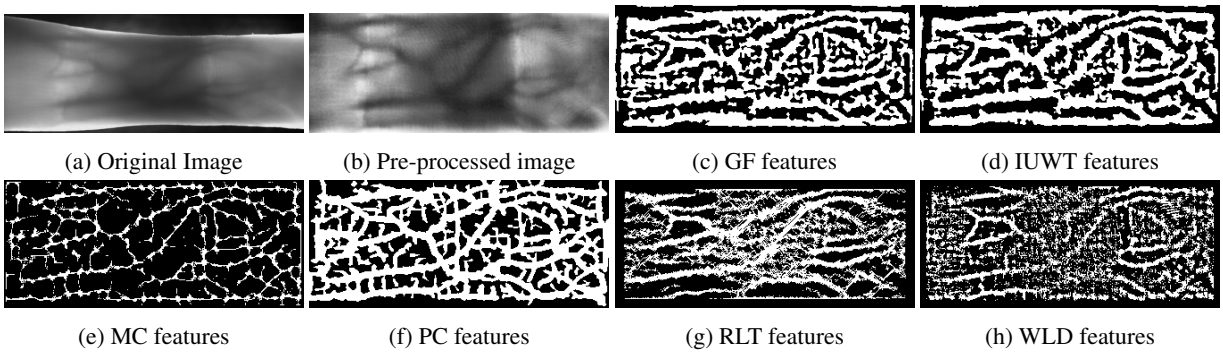

(a) Original Image (b) Pre-processed image (c) GF features (d) IUWT features

(e) MC features (f) PC features (g) RLT features (h) WLD features

Figure 1: Original UTFVP finger vein image, pre-processed image, and extracted binary features.

2.2. Feature Extraction

Feature representations are derived from pre-processed images, with the aim of extracting discriminative information from them. In order to perform a proper comparison of the three considered BTP approaches in terms of achievable recognition performance, we have benchmarked different feature extraction algorithms. All of them generate binary templates containing geometric information related to the shape or topological structure of the observed vein patterns. The methods considered are the following:

- **Gabor Filtering** (GF) [23]: inspired by the human visual system with its multi-channel processing of visual information, this approach processes vein data through a bank of kernels to obtain distinct filtered images. These images are subsequently fused into a single representation to generate the desired binary template by thresholding the sum of the individual outputs from each filter;

- **Isotropic Undecimated Wavelet Transform** (IUWT) [11] is a redundant wavelet transform whose coefficients encode information corresponding to different spatial scales. Levels 2 and 3 of such transform exhibit the best contrast for the blood vessels, and are therefore used to create the sought binary template, by thresholding the obtained images and post-processing the results using morphological operations to remove residual noise;

- **Maximum Curvature** (MC) [29] extracts the lines corresponding to the central part of the veins, thereby producing templates robust to varying vein widths. The vein locations are first estimated using curvature information extracted from the input image, then refined connecting dots through a filtering operation, before generating the final binary output comparing the obtained image with the corresponding median value;

- **Principal Curvature** (PC) [5] computes the gradient field of the input image, and then applies filtering to remove noise and smooth the output. The binary template is finally extracted from the principal curvature information, which is obtained from the eigenvalues of the Hessian matrix at each pixel;

- **Repeated Line Tracking** (RLT) [28] estimates a statistical likelihood of each pixel belonging to a blood vessel. This estimation is carried out over the input image in an iterative manner, where veins are tracked starting from a random point, and with a final binarisation step;
- **Wide Line Detector** (WLD) [16] is an adaptive thresholding technique, comparing each pixel with its neighbourhood to determine which ones should represent veins in the final binary template.

Figure 1 shows the binary templates extracted from a single sample, which are subsequently used as input for the considered template protection schemes. The publicly available open-source PLUS-OpenVein Toolkit[1] has been used to process finger vein images.

2.3. Comparison

The binary templates generated from two vein images can be compared in terms of their correlation [29]. Since the input images are only coarsely aligned to each other, we compute the correlation between the probe image $I(x, y)$ and several rotated and shifted versions of the reference image $R(x, y)$:

$$N_m(s,t,\theta) = \sum_{y=0}^{h-2c_h-1} \sum_{x=0}^{w-2c_w-1} I(s+x, t+y) \cdot Rot_\theta(R)(c_w+x, c_h+y)$$

where s and t are the possible shift values in horizontal and vertical directions; h and w are the image height and width; c_h and c_w are the number of pixels to shift in vertical and horizontal direction, $Rot_\theta(R)$ refers to the image R rotated by an angle θ, and $N_m(s,t,\theta)$ is the corresponding correlation value. The maximum value of the correlation over all shifts and rotations is then defined as:

$$N_{m_{max}} = \max_{0\le s<2c_w, 0\le t<2c_h, -c_r\le\theta\le c_r} N_m(s,t,\theta)$$

Finally, the maximum value of the correlation is normalised and used as comparison score [29]:

$$score = \frac{N_{m_{max}}}{\sum_{y=t_0}^{t_0+h-2c_h-1} \sum_{x=s_0}^{s_0+w-2c_w-1} I(x,y) + \sum_{y=c_h}^{h-2c_h-1} \sum_{x=c_w}^{w-2c_w-1} Rot_{\theta_0}(R)(x,y)}$$

where s_0, t_o and θ_0 are the indices of $N_{m_{max}}$ in the correlation matrix $N_m(s,t,\theta)$. The resulting score values are in the range $[0, 0.5]$.

In order to compare templates protected using block remapping and block warping, the above mentioned formula is utilised. In contrast, an averaged Hamming distance as described in Section 3.3 is used for Bloom-filter templates.

[1] http://wavelab.at/sources/OpenVein-Toolkit/

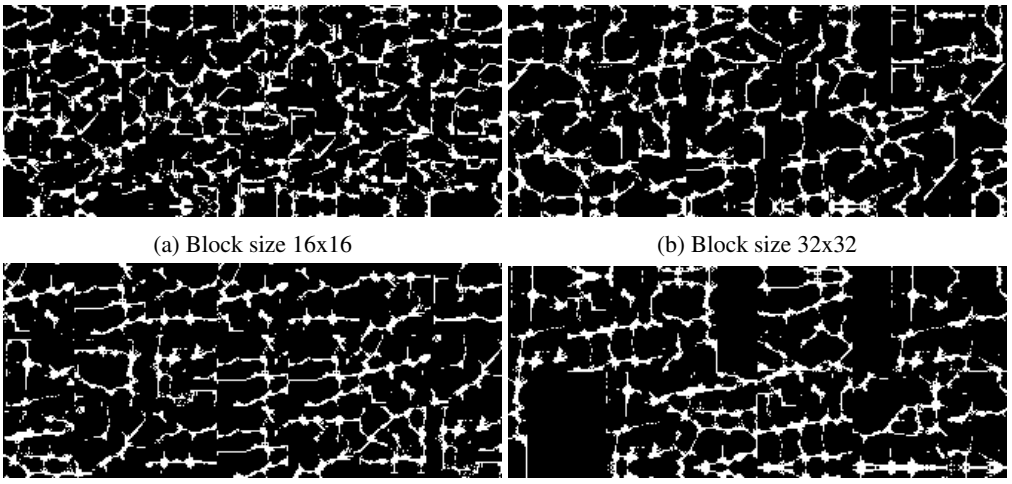

(a) Block size 16x16 (b) Block size 32x32

(c) Block size 48x48 (d) Block size 64x64

Figure 2: Block remapping example with different block sizes using the same MC feature image as in Fig. 1e

3. Finger Vein Cancelable Biometrics

The non-invertible transforms used to generate cancelable templates from the finger vein binary representations described in Section 2 are detailed in this section. All the employed transforms rely on a single system-specific key; i.e., the same key can be used for each user in the system, not requiring any specific handling or secure storage of user-specific keys.

3.1. Block Remapping

A fixed-size region of $N \times W$ pixels, aligned to the center of the finger area, is extracted from the binary template. The selected region is then divided into B_T square blocks of $B \times B$ pixels, out of which a subset of B_C blocks are randomly selected. The chosen blocks are remapped according to a system-dependent pre-defined key [36]. The obtained distorted templates can be compared against each other in the transformed domain as described in Section 2.3. Examples of cancelable templates generated with this scheme from MC features are presented in Figure 2.

It should be noted that not all the original blocks will appear in the remapped template, thereby granting non-invertibility. The selected blocks can be repeated multiple times in the transformed template, in order to obtain an image whose size is the same as the input one. The decisions about which blocks to consider, and their positions in the remapped template, depend on the employed key. Furthermore, different blocks may not contain the same amount of vein information. Thus, the recognition performance can be affected by such selection. Specifically, due to illumination issues, blocks extracted from the outer part of the finger generally contain less vein information compared to the ones belonging to the central area. We therefore perform a key-space reduction considering only the selection of blocks belonging to the central area of the finger. Besides the employed key, also the used block size influences

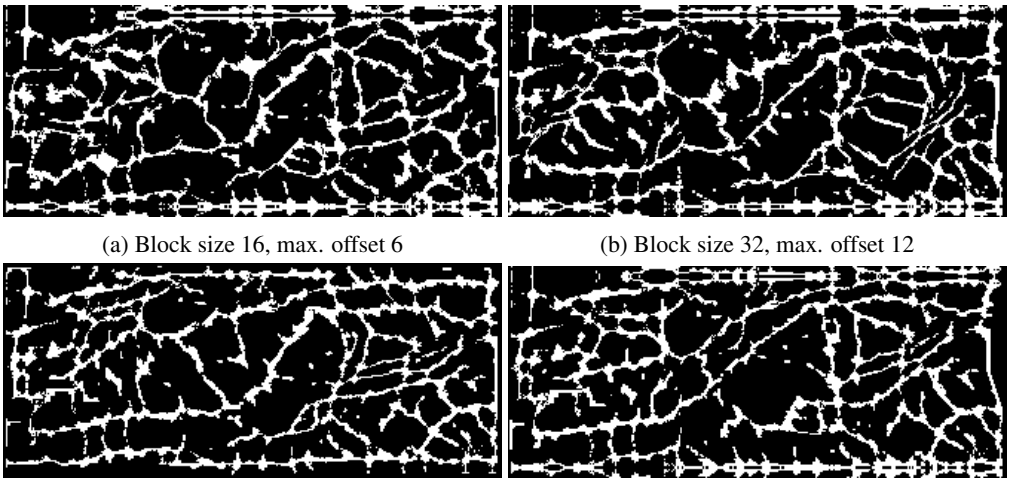

(a) Block size 16, max. offset 6 (b) Block size 32, max. offset 12

(c) Block size 48, max. offset 18 (d) Block size 64, max. offset 24

Figure 3: Block warping example with different parameters using the same MC feature image as in Fig. 1e

the scheme performance. Actually, there is a trade-off between security and recognition performance related to the block size: a smaller block size leads to higher security in terms of irreversibility and unlinkability, yet decreases the recognition performance, whereas a bigger block size slightly affects recognition performance, yet lowers the provided security [32]. The impact of the aforementioned parameters on recognition performance, as well as on unlinkability, is evaluated in Section 5.2.1.

3.2. Image Warping

Following the mesh warping algorithm [46], a grid is laid over the binary template, and its vertices are offset by amounts linked to the key defining the transformation. Each row and column of the template is then transformed by carrying out a miniaturisation or an expansion based on the distorted grid. Miniaturisation is performed with the help of a box filter, while linear interpolation is used for expansion. Similarly to block remapping, there is a trade-off between security and recognition performance. The more distorted the warped grid is compared to the regular grid, the higher the amount of interpolation applied, and the more secure the transformation is. On the other hand, the recognition accuracy decreases with a higher level of distortion and vice versa. The distorted templates, whose examples are shown in Figure 3, can still be compared as described in Section 2.3.

3.3. Bloom Filters

Bloom filters are here employed to generate cancelable templates according to the scheme introduced in [9], which improves the unlinkability properties of the original approach proposed in [37]. In more detail, the following three steps are performed:

1. *feature extraction and encoding*: the original two-dimensional binary template is divided into $nBlocks$ blocks,

each having $nBits \times nWords$ bits;

2. *structure-preserving feature re-arrangement*: this step, proposed in [9], is performed to accomplish the desired unlinkability of the produced templates, by randomly spreading the available information into different blocks, while trying to preserve the recognition accuracy. To this end, the $nBlocks$ blocks of the unprotected template are re-grouped into $nBlocksX$ sets, each consisting of $nBlocksY$ blocks ($nBlocks = nBlocksX \times nBlocksY$). Then, the rows of the vertical concatenation of the $nBlocksY$ blocks are permuted within each set, according to a predefined key. This way, it is not possible to exploit the Hamming weights of corresponding blocks to link templates enrolled in different systems, as proposed in [13];

3. *Bloom filter computation*: protected templates are extracted by computing one Bloom filter $\mathbf{b}$ from each of the $nBlocks$ blocks, such that the final protected template $\mathbf{C}$ consists of $nBlocks$ Bloom filters of size 2^{nBits}: $\mathbf{C} = \left\{\mathbf{b}^{(1)}, \ldots, \mathbf{b}^{(nBlocks)}\right\}$. In order to map one block to a Bloom filter, the $nWords$ columns of each block are translated to their decimal value, and the corresponding indices are set to one in the Bloom filter. These steps add irreversibility to the templates, since the number of occurrences of each column, and their positions within the block, are lost.

The final comparison score s between a probe $\mathbf{C}_q$ and a reference $\mathbf{C}_r$ protected templates is defined as the average Bloom-filter-based dissimilarity score:

$$S(\mathbf{C}_q, \mathbf{C}_r) = \frac{1}{nBlocks} \sum_{i=1}^{nBlocks} \frac{\mathrm{HD}(\mathbf{b}_q^{(i)}, \mathbf{b}_r^{(i)})}{|\mathbf{b}_q^{(i)}| + |\mathbf{b}_r^{(i)}|} \tag{1}$$

where $|\mathbf{b}|$ is the number of bits set to 1 within a Bloom filter $\mathbf{b}$, and $\mathrm{HD}(\mathbf{b}_q^{(i)}, \mathbf{b}_r^{(i)})$ the Hamming distance between two filters.

3.4. Pre-alignment for Template Protection

Misalignment of two templates, in terms of shifts and planar rotations, does not only cause issues during the comparison, but even more severe problems for all block-based cancelable biometrics schemes. While planar rotations and vertical shifts can be ruled out easily for finger veins, as described in Section 2.1, dealing with horizontal shifts is not as straightforward. Horizontal shifts are usually compensated during comparison by shifting one of the templates. However, this strategy cannot be used for protected templates, especially if the shifts become larger than the employed block size: feature information in different blocks is treated differently, thereby leading to a dissimilar output template. Thus, a suitable pre-alignment prior to the application of the block-based cancelable scheme is needed.

We propose a pre-alignment strategy in which all templates originating from the same finger are first registered

against each other, with the help of a correlation-based approach derived from [29]. Specifically, the employed scheme provides the relative position of the two templates to each other. This position encodes the shifts needed during comparison to achieve the highest score, i.e., the best possible alignment. Hence, if the reference template is shifted according to that information, an alignment/registration of the templates can be done. This alignment is done only if the comparison score is above a pre-defined threshold, in order to avoid ambiguous outcomes and it requires the unprotected templates to be available. Its benefits are shown in the recognition performance evaluation of the employed cancelable biometrics schemes in Section 5.2.3.

4. Security Analysis

The security analysis of the employed schemes is conducted in terms of unlinkability and irreversibility, in compliance with the ISO/IEC 24745 standard on biometric information protection [17]. A renewability analysis of block remapping and warping is done in [32], and in [10] for Bloom Filters. To quantify unlinkability, a general approach based on comparison scores is here employed, as outlined in Section 4.1. Regarding irreversibility, beyond general considerations about the considered BTP schemes, an attack against the block remapping approach, based on automated square jigsaw puzzle solvers, is presented and evaluated in Section 4.2, where also the irreversibility of block warping and Bloom filter approaches is briefly discussed. Furthermore, the possibility of recovering the original biometric information by exploiting the availability of multiple cancelable templates generated from the same biometric data is discussed in Section 4.2.1.

4.1. Unlinkability

Template unlinkability is evaluated according to the protocol proposed in [8] and publicly available [2]. Two protected templates $\mathbf{T}_1$ and $\mathbf{T}_2$ generated from the same biometric sample are defined as linkable if an attacker can determine that they were extracted from mated instances, and hence conceal a unique identity.

To accomplish such goal, the attacker compares the two protected templates by computing a linkage score $s = LS\left(\mathbf{T}_1, \mathbf{T}_2\right)$, upon which a decision regarding whether the considered templates actually stem from mated instances is taken. This linkage score is obtained by comparing the protected templates through a distance metric. In case of Bloom filters, the Hamming distance can be used, while the same comparison approach employed for unprotected templates can be exploited for the block remapping and warping. Following Kerckhoffs's principle, it is assumed that the attacker knows how the system works and, in particular, the *mated-instance* and *non-mated-instance* score distributions generated by comparing protected templates. These distributions can be quantitatively compared by means of two different measures:

[2] https://github.com/dasec/unlinkability-metric

- a local measure $\mathrm{D}_{\leftrightarrow}(s)$, evaluating the linkability of templates on a score-wise basis. A measure $\mathrm{D}_{\leftrightarrow}\left(s_1\right) = 1$ for a specific linkage score s_1 means that an attacker is able to link the considered templates to the same instance with full certainty. On the other hand, $\mathrm{D}_{\leftrightarrow}\left(s_0\right) = 0$ should be interpreted as full unlinkability for templates giving linkage score s_0. Intermediate values of $\mathrm{D}_{\leftrightarrow}(s)$ report an increasing degree of linkability;

- a global measure $\mathrm{D}_{\leftrightarrow}^{sys}$, giving an overall evaluation of the whole BTP scheme unlinkability. A system with $\mathrm{D}_{\leftrightarrow}^{sys} = 1$ should be fully linkable, meaning that mated-instance and non-mated-instance score distributions have no overlap, and local measures $\mathrm{D}_{\leftrightarrow}(s) = 1$ for linkage scores computed from any pair of mated samples. Similarly, $\mathrm{D}_{\leftrightarrow}^{sys} = 0$ means that the system is fully unlinkable, with mated and non-mated score distributions completely overlapping. All intermediate values of $\mathrm{D}_{\leftrightarrow}^{sys}$ report a decreasing degree of unlinkability.

As detailed in [8], given a linkage score s, the local measure $\mathrm{D}_{\leftrightarrow}(s)$ indicates whether it is more likely that the two considered templates stem from mated instances, whose probability is $p\left(H_M|s\right)$ for hypothesis H_M, than from non-mated instances, characterized by probability $p\left(H_{NM}|s\right)$ for hypothesis H_{NM}. Therefore, $\mathrm{D}_{\leftrightarrow}(s)$ can be expressed as the difference of conditional probabilities for each hypothesis:

$$D_{\leftrightarrow}(s) = p\left(H_M|s\right) - p\left(H_{NM}|s\right). \tag{2}$$

The unknown conditional probabilities can be expressed through the known probabilities of obtaining s given templates belonging to mated or non-mated samples, that is, $p\left(s|H_M\right)$ and $p\left(s|H_{NM}\right)$. We can rewrite eq. (2) in terms of the likelihood ratio between them, $LR(s) = p\left(s|H_M\right)/p\left(s|H_{NM}\right)$ as:

$$\mathrm{D}_{\leftrightarrow}(s) = \begin{cases} 0 & \text{if } LR(s)\cdot\omega \leq 1 \\ 2\frac{LR(s)\cdot\omega}{1+LR(s)\cdot\omega} - 1 & \text{if } LR(s)\cdot\omega > 1 \end{cases} \tag{3}$$

where $\omega = p\left(H_M\right)/p\left(H_{NM}\right)$ is the ratio between the prior probabilities of mated and non-mated samples distributions. This latter ratio can be assumed to be known for operating systems with registered mated and non-mated access attempts [8], or can be set to $\omega = 1$ as in the present analysis.

The global linkability measure is instead computed measuring how likely it is to get a linkage score stemming from the mated samples distribution, being then defined in [8] as:

$$\mathrm{D}_{\leftrightarrow}^{sys} = \int_{s_{min}}^{s_{max}} p\left(s|H_m\right)\cdot \mathrm{D}_{\leftrightarrow}(s)\,\mathrm{d}s. \tag{4}$$

4.2. Irreversibility

Measuring the irreversibility of a BTP scheme means evaluating the amount of information regarding the original biometric template or sample which the protected one leaks. We can consider two different scenarios: in the first one, the template protection key is known to the attacker, i.e., it has been compromised, while in the second one the key is not known to the attacker.

Regarding the block remapping scheme, irreversibility has an upper bound given by the number of blocks selected from the original template to produce the protected one. The original template could be partially reconstructed from its protected version through a brute-force attack, where all possible block permutations are tested until the correct block order is found. However, in order to conduct such an attack without the knowledge of the original content, an indicator is needed to establish whether a block is set to the correct position or not. This could be done by comparing the border pixels of the available blocks, and searching for the best match among possible combinations. Such task essentially looks like solving a square jigsaw puzzle, an issue for which several automated procedures can be found in the literature [31, 4]. Since the protected templates do not contain all blocks of the original representations, an appropriate reconstruction approach should be able to deal with missing parts. This can be done using the greedy placement strategy and the prediction-based dissimilarity metrics proposed in [31], for which a public implementation is available[3]. This approach does not require any prior knowledge about the original data, and is able to handle puzzles with missing pieces, pieces of unknown orientation, and unknown overall size. The performance of a square jigsaw puzzle solver can be measured according to either global and local metrics [4]. The former ones compare original and reconstructed contents quantifying the number of blocks at the correct position. The latter ones focus on clusters of blocks, rating either the biggest correct block cluster or the number of correct block pairs, that is, blocks with at least one correct neighbour. Since in our scenario the amount of information leakage does not depend on the absolute block positions but on the continuity of the vein lines (local clusters), both local metrics are employed in the evaluation in Section 5.3.2.

Regarding the block warping scheme, it is possible to derive the key if the key is not known to the attacker, or at least obtain some hints about it, by analysing the interpolation artifacts using image forensic methods [2]. Depending on the strength of the applied warping, the key can be restored with a certain probability. Nevertheless, the applied mesh warping transformation can be considered irreversible under both scenarios (known or unknown key), since interpolation strategies are applied, thereby leading to a loss of information. Thus, it is not possible to completely recover the original data even if the warping parameters are known. Furthermore, the level of irreversibility is higher if miniaturisation is applied, due to overlay effects [12].

Finally, in a full disclosure model where the attacker knows both the stored templates and the employed key within

[3] https://github.com/ZaydH/sjsu_thesis

the Bloom filter scheme [9], the number $nSeq$ of possible $nBits \times nWords$ original binary representations which could be derived from a given Bloom filter, corresponding to one of the $nBlocks$ blocks, is given by

$$nSeq = \sum_{i=1}^{|\mathbf{b}|} (-1)^{|\mathbf{b}|-i} \binom{|\mathbf{b}|}{i} i^{nWords}, \tag{5}$$

being $|\mathbf{b}|$ the Hamming weight of the Bloom filter associated with the considered block. This number rapidly reaches high values, greater than 10^{20}, even for small $|\mathbf{b}|$ weights. Such large search space makes it hard to conduct brute-force attacks, leading to the desired irreversibility. To the best of our knowledge, for cancelable biometrics so far there is no agnostic attack addressing the irreversibility of Bloom filters. It is yet worth pointing out that there are studies evaluating the robustness of Bloom-filter-based privacy-preserving record linkage (PPRL) applications, such as those involved in the protection of medical records [39], against attacks aimed at recovering the encrypted information [43]. However, the attacks considered in the aforementioned scenario are commonly based on specific statistical knowledge about the data to be protected, usually consisting of people names, and on the exploitation of dictionaries and language models to reduce the space of possible binary patterns which could be generated. In the context of the protection of biometric templates, such additional information is typically not available, and the involved bit distributions often approximates uniformity. It is worth mentioning that strategies based on simple operations applied to the employed Bloom filters can be adopted to notably increase their robustness to statistical attacks [35].

4.2.1. Record Multiplicity Attacks

Besides trying to retrieve the original biometric information from a stolen protected template, an attacker could also resort to more sophisticated approaches, such as those involving more than a single cancelable biometrics derived from the same original sample. Recovery approaches relying on multiple protected representations, usually taken from different databases, are commonly referred to as record multiplicity attacks (RMAs) [26], correlation or coalition attacks [30].

In order to launch such attacks, it is required to know which protected templates, enrolled in distinct applications, have been derived from the same original biometric data. Unlinkability across different representations obtained from the same biometric instance is a property not only useful to preserve the privacy of the interested subjects, but also to limit the feasibility of RMAs, making it hard for an attacker to know which templates are linked to the same identity. As a consequence, evaluating the unlinkability of a BTP scheme as discussed in Section 4.1 already provides valuable information regarding its robustness to RMAs.

Nonetheless, it has to be observed that an attacker, besides resorting to cross-matching scores, can exploit other strategies to link templates related to the same original biometric instance, relying for instance on information transferred on side channels. An in-depth evaluation of the threats associated to RMAs should be therefore addressed even

if proper template unlinkability is provided. Regarding the BTP schemes considered in this paper, the robustness of the block remapping approach against RMAs has been investigated in [19], where it has been shown that the irreversibility of such method is greatly diminished if multiple templates, derived from the same original iris instance, are available to an attacker. The evaluations reported in [19] are applicable also to the vein templates considered here, thus making the block remapping BTP vulnerable to RMAs.

Similarly, it has been shown that it is possible to exploit the knowledge derived from the availability of more than a single cancelable minutiae-based fingerprint template to reduce the non-invertibility of the employed many-to-one block warping functions [34]. Applying the approaches in [34] to the binary vein representations here considered would imply an increased computational complexity, and require the availability of a larger number of compromised templates due to the greater amount of points to which the inverse warping has to be applied, with respect to a scenario involving minutiae templates. Nonetheless, a RMA against a block warping BTP would be unquestionably successful, in the sense that it would entail a gain on the amount of information extracted from the combination of protected templates, with respect to having access to a single template.

Finally, to the best of our knowledge there is no study in literature evaluating the effectiveness of RMAs against Bloom-filter-based protection schemes. Although performing a rigorous and in-depth analysis about the robustness of Bloom filters to possible RMAs is out of the scope of the present paper, it can be observed that, even when resorting to the original proposal in [37] instead of the improved one here adopted, it would be hard to jointly exploit the information within two distinct Bloom-filter-based protected templates, obtained from the same original data. In fact, it is worth remarking that the security of Bloom filters as in [37] does not rely on the exploitation of a user-specific secret key, which is instead employed only to provide renewability. Conversely, it is based on the many-to-one mapping of multiple columns of a block to the same bit of the filter, and also on the loss of knowledge about which column has contributed to a particular bit in the generated filter. Therefore, two Bloom filters generated from the same biometric data according to the approach in [37] would only leak information associated to the Hamming weight $|\mathbf{b}|$, which would be the same for the two filters, without the possibility of combining any further information, thus leaving the search space of an attack relying on multiple templates expressed by the term in eq. (5). Furthermore, the enhanced implementation adopted here makes it even more difficult to efficiently combine the information leaked from two distinct protected templates, since in this case they are generated after distinct random permutations, which scramble the inner structures of the considered blocks. Actually, performing random permutations in a BTP scheme typically improves the achievable irreversibility, as proven for helper-data-based biometric cryptosystems in [20]. Therefore, the search space would be simply limited by the permutation generating blocks with the lowest Hamming weights $|\mathbf{b}|$. However, it would be still computationally hard to conduct a successful attack, since eq. (5) returns high values even for low $|\mathbf{b}|$ weights. It could be therefore argued that, differently from using block remapping or

block warping BTP schemes, the adoption of a BTP scheme, relying on Bloom filters, makes the created cancelable biometrics able to withstand RMAs.

5. Experimental Evaluation

The tests experimental setup is detailed in Section 5.1. The employed recognition performance evaluation protocol is then outlined in Section 5.2. The tests regarding the performed security analysis of the considered template protection schemes are presented in Section 5.3. Finally, a discussion summarising the obtained results is reported in Section 5.4.

5.1. Experimental Setup

5.1.1. Finger Vein Dataset

Two publicly available finger vein datasets are used in our experiments: UTFVP and SDUMLA-HMT.

The UTFVP (University of Twente Finger Vascular Pattern Database) [41] dataset contains 1440 images in total, captured from 60 subjects during two distinct recording session, with 6 fingers per subject (index, middle and ring finger) and 4 images per finger. The images have a resolution of 672×380 pixels. The width of the visible vein lines inside the images is between $4 - 20$ pixels. The binary templates extracted from the pre-processed images have a size of $N \times W = 336 \times 142$ pixels.

The SDUMLA-HMT [47] finger vein dataset is composed of 3816 images in total, with each of 6 fingers (index, middle and ring finger of both hands) captured for 6 times from 106 subjects. The images are recorded as 8-bit greyscale, and stored in BMP format with a resolution of 320×240 pixels.

The images of both the databases are processed in the same way: at first the images are pre-processed, as described in Section 2, extracting the finger vein region and performing an image enhancement. Then the respective features (GF, IUWT, MC, PC, RLT or WLD) are extracted from the pre-processed images, resulting in binary feature images of size 336×142 pixel and 160×64 pixel for the UTFVP and SDUMLA, respectively. Afterwards, the considered cancelable biometrics approaches (*block remapping*, *image warping*, and *Bloom filter*) are applied to the aforementioned binary templates. Figure 4 shows some example pairs of protected templates from different sessions for each template protection scheme.

5.1.2. Set-up of the Employed Approaches

The parameters used for the three considered BTP schemes are described in the following:

- **block remapping**: the input feature template is divided into square blocks of size $B \times B$ pixels, with $B = \{16, 32, 48, 64\}$ in the performed tests. Cropping is performed in case the template dimensions are not multiple

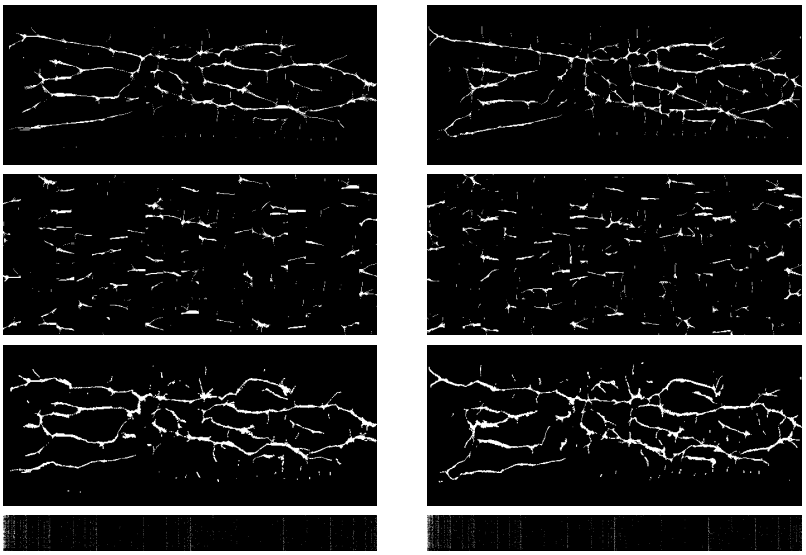

Figure 4: Pairs of protected samples (for the same instance) from two distinct sessions (first and second column). First row: original templates (MC), second row: block rempapping (32x32), third row: block warping (32-12), fourth row: Bloom filters.

of B. As for the percentage of original blocks to be included in the transformed template, we have empirically evaluated that keeping it at 75% leads to the best trade-off between recognition performance and privacy. The selected blocks are rearranged according to a system-dependent key in the remapped template, whose size is same as the original input;

- **block warping**: a rectangular grid of $G \times G$ pixels is laid on the template, and the maximum vertices offset O the transformation can apply is set. Such offset should be smaller than half the block size to generate usable outputs. Cropping is performed in case the size of the original template does not allow using a grid with all equal elements. In particular, we have tested the pairs $\{G = 16, O = 6\}$, $\{G = 32, O = 12\}$, $\{G = 48, O = 18\}$, and $\{G = 64, O = 24\}$;

- **bloom filters**: the settings of the Bloom-filter-based protection scheme are determined as in [10]. Specifically, for the UTFVP database $nBits$ is empirically set to 10 to balance irreversibility and recognition performance, while $nWords$ is selected within the allowed range to 48 in order to maintain proper recognition capabilities. The number of blocks in horizontal direction $nBlocksX$ is set to 7, while $nBlocksY = 14$ for the vertical direction, with $nBlocks = nBlocksX \cdot nBlocksY$. For the SDUMLA database, we use $nWords = 22$, $nBits = 4$, $nBlocksX = 7$ and $nBlocksY = 16$. In addition, 4 different key permutations are used to add unlinkability to the Bloom-filter templates.

Table 1
Baseline results on the UTFVP and the SDUMLA-HMT dataset.

Dataset		GF	IUWT	MC	PC	RLT	WLD
UTFVP	EER	0.77%	0.77%	**0.36%**	0.57%	2.1%	0.72%
	FMR1000	1.33%	1.28%	**0.51%**	0.92%	4.05%	1.28%
	ZeroFMR	4.26%	4.41%	**2.15%**	4.21%	11.9%	6.1%
SDUMLA-HMT	EER	6.55%	5.25%	**3.68%**	4.95%	11.6%	5.66%
	FMR1000	12.8%	9.75%	**5.95%**	7.76%	22.7%	9.91%
	ZeroFMR	**23.8%**	86.1%	86.7%	96.7%	74.0%	90.5%

5.2. Recognition Performance Evaluation

The system recognition performance is measured in terms of the false non-match rate (FNMR) and corresponding false match rate (FMR). The variance of the obtained equal error rate (EER) is used as an indicator for the dependence of the recognition performance on the employed transformations key parameter. In addition, the FMR1000 (the lowest FNMR for FMR = 0.1%) and the ZeroFMR (the lowest FNMR for FMR = 0%) are used to quantify the recognition performance. Section 5.2.1 presents the results obtained for baseline unprotected systems, whereas Sections 5.2.2 and 5.2.3 outline the achievements of the proposed cancelable biometric approaches, without and with the proposed pre-alignment procedure, respectively.

5.2.1. Baseline Recognition Performance Results

The baseline results in terms of EER, FMR1000 and ZeroFMR for the six employed feature representations are listed in Table 1 for both the UTFVP and the SDUMLA-HMT finger vein databases. On the UTFVP dataset, the reported values show that MC performs best in terms of all three performance indicators, followed by PC, IUWT, GF, and WLD, while RLT performs worst. On the SDUMLA dataset, MC performs best in terms of EER and FMR1000 whereas GF performs best in terms of ZeroFMR.

5.2.2. Cancelable Schemes Recognition Performance Results

Table 2 and Table 3 report the performance of the considered cancelable biometrics schemes, in terms of mean EER with the corresponding 95% confidence intervals on the UTFVP and SDUMLA-HMT datasets, respectively. The detection error trade-off (DET) curves showing the aforementioned results are reported in Figure 5 and 6 for the UTFVP and the SDUMLA-HMT datasets, respectively. For the UTFVP dataset, Figure 7 shows the impact of parameter selection for block remapping and block warping for the MC feature representation. For block remapping, bigger block sizes are preferable for recognition purposes, with MC performing overall best. The same holds for block warping, where bigger block sizes lead to a better performance despite the higher maximum offset, where MC performs overall best for a block size of 64 and a maximum offset of 24. The best performance using Bloom filters is

Table 2
Recognition performance results in terms of EER and 95% confidence intervals for cancelable biometrics schemes applied on the UTFVP database, using 10 different transformation keys for each template.

Scheme	Parameters	GF	IUWT	MC	PC	RLT	WLD
Block Remapping	16x16	14.9%±0.74	15.5%±0.75	**12.5%±0.7**	12.6%±0.68	17.2%±0.77	13.2%±0.7
	32x32	6.58%±0.66	6.99%±0.66	5.26%±0.58	**5.05%±0.57**	8.33%±0.68	5.47%±0.58
	48x48	11.8%±0.65	11.7%±0.63	**6.5%±0.51**	6.9%±0.54	9.72%±0.64	6.87%±0.55
	64x64	7.36%±0.55	7.69%±0.56	**3.96%±0.41**	4.91%±0.47	7.64%±0.55	5.2%±0.69
Block Warping	16 - 6	3.78%±0.39	2.08%±0.32	**1.29%±0.23**	2.35%±0.31	3.15%±0.38	3.66%±0.38
	32 - 12	3.38%±0.43	2.59%±0.39	**1.25%±0.31**	2.18%±0.36	3.66%±0.42	2.6%±0.39
	48 - 18	2.68%±0.36	2.59%±0.34	**1.33%±0.27**	1.67%±0.29	3.2%±0.39	1.76%±0.33
	64 - 24	1.57%±0.34	1.25%±0.31	**0.73%±0.23**	1.29%±0.27	2.04%±0.35	1.11%±0.31
Bloom Filters	48/10/7/14	16.6%±0.79	13.5%±0.73	18.3%±0.8	**11.4%±0.67**	14.7%±0.74	14.2%±0.73

Table 3
Recognition performance results in terms of EER and 95% confidence intervals for cancelable biometrics schemes applied on the SDUMLA-HMT database, using 10 different transformation keys for each template.

Scheme	Parameters	GF	IUWT	MC	PC	RLT	WLD
Block Remapping	16x16	31.7%±0.47	24.3%±0.45	**23.1%±0.44**	23.6%±0.44	25.5%±0.46	24.7%±0.46
	32x32	30.3%±0.47	22.7%±0.44	**21.5%±0.43**	22.4%±0.43	26.1%±0.45	23.9%±0.45
	48x48	28.2%±0.46	17.4%±0.42	**12.9%±0.41**	16.9%±0.42	20.3%±0.44	16.4%±0.42
	64x64	17.0%±0.42	4.9%±0.25	**3.45%±0.24**	4.85%±0.25	10.4%±0.33	5.52%±0.26
Block Warping	16 - 6	17.3%±0.38	8.5%±0.32	**7.85%±0.33**	8.12%±0.32	14.0%±0.37	9.93%±0.34
	32 - 12	16.5%±0.39	6.49%±0.27	**5.21%±0.26**	6.51%±0.27	12.4%±0.34	7.36%±0.29
	48 - 18	22.6%±0.44	18.2%±0.4	16.7%±0.32	17.4%±0.39	**15.8%±0.38**	18.0%±0.4
	64 - 24	14.2%±0.38	5.34%±0.23	**3.76%±0.2**	5.02%±0.22	11.5%±0.33	5.81%±0.24
Bloom Filters	22/4/7/16	29.2%±0.47	23.7%±0.44	22.8%±0.43	24.6%±0.44	23.3%±0.43	**21.0%±0.42**

achieved for the PC-based features, yet with overall results far worse than those achieved with the block remapping and the block warping BTP schemes.

On the SDUMLA dataset, the MC features achieves the best results for most combinations, except for block warping with a block size of 48 and a maximum offset of 18, and for Bloom filters. As on the UTFVP dataset, the same trend regarding block sizes for block remapping, and block size/offset combinations for block warping, is visible: higher block sizes lead to an improved recognition performance in terms of EER. In general, again block warping achieves a much higher performance than block remapping, while the Bloom filter approach is ranked last in terms of recognition performance. The overall best performance on the SDUMLA dataset is achieved using block remapping with a block size of 64×64 for the MC features.

In general, it can be observed that template protection significantly degrades the achievable recognition perfor-

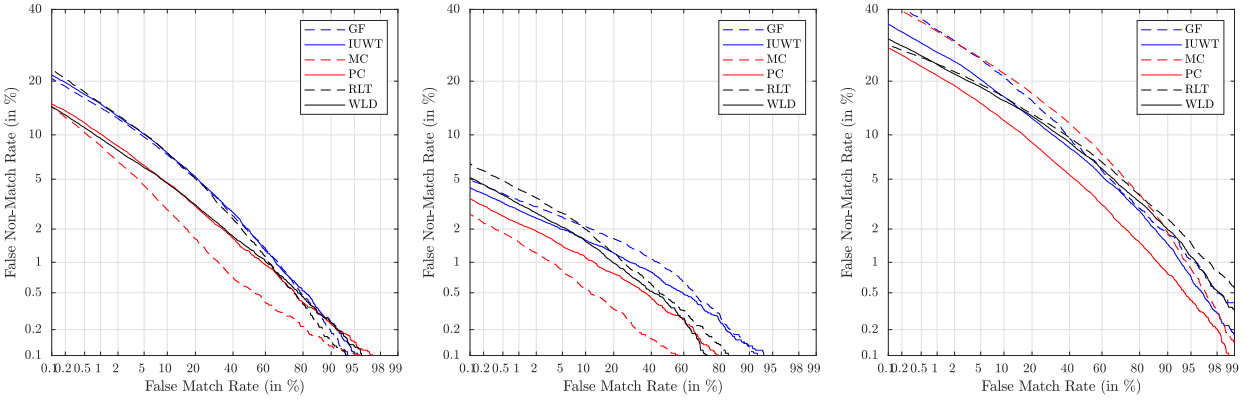

(a) Block Remapping 64x64 (b) Block Warping 64 - 24 (c) Bloom Filter

Figure 5: DET curves for different features and BTP schemes: (a) block remapping, (b) block warping and (c) Bloom Filters on the UTFVP database.

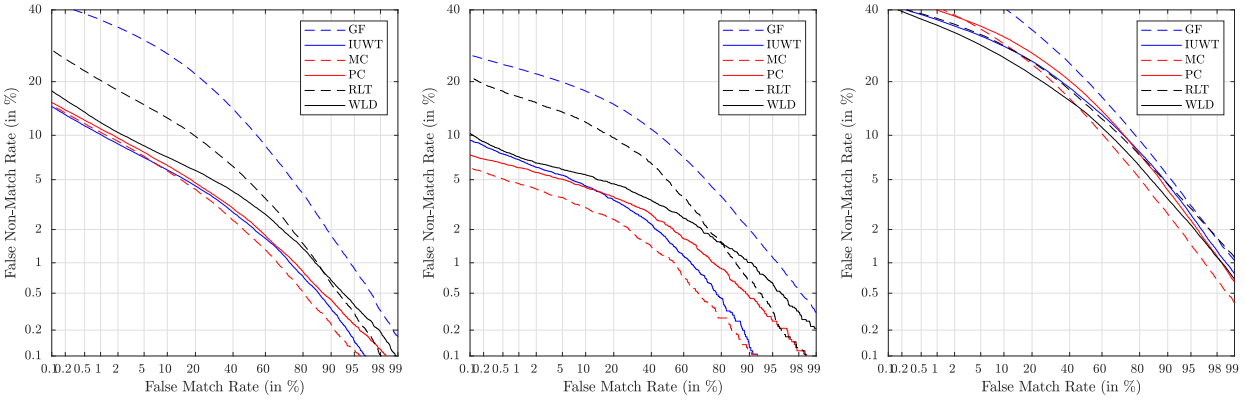

(a) Block Remapping 64x64 (b) Block Warping 64 - 24 (c) Bloom Filter

Figure 6: DET curves for all six feature types and the different template protection schemes: (a) block remapping, (b) block warping and (c) Bloom Filters on the SDUMLA-HMT database.

mance, while the employed transformation keys introduce only limited variability in the obtained results (low variance within the 95% confidence intervals). The main reasons for the performance degradation are the shifts and rotations present in the input samples. All the tested template protection schemes are sensitive to misaligned input templates. To further motivate our proposed pre-alignment strategy, for which results are presented in the following, Figure 8 shows an example of non-aligned and pre-aligned protected templates (using block remapping 32×32) from the same instance, including the obtained comparison scores. It is clearly visible that the pre-alignment greatly improves the comparison score.

5.2.3. Recognition Performance Results with Pre-Alignment

The effectiveness of the alignment approach proposed in Section 3.4 for the employed cancelable schemes is confirmed by the results given in Table 4 and Table 5 for the UTFVP and the SDUMLA dataset, respectively. Considering

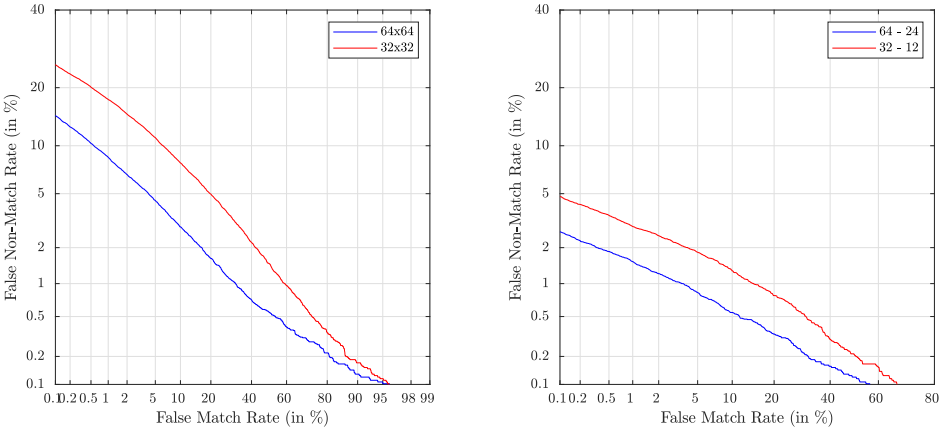

(a) Block Remapping (b) Block Warping

Figure 7: DET curves for different parameters in BTP schemes applied to MC features on the UTFVP dataset: (a) block remapping, (b) block warping.

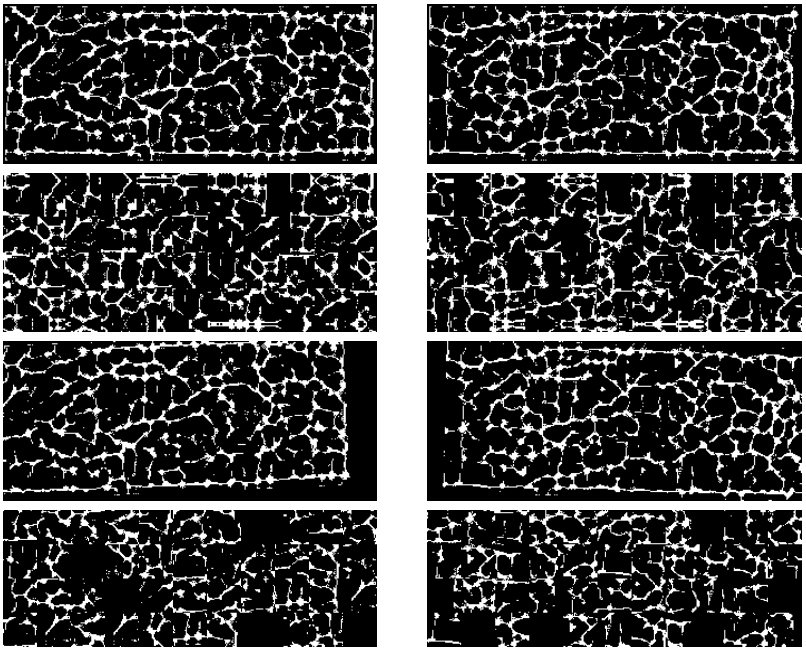

Figure 8: Example of the pre-alignment. First row: original input templates - comparison score 0.1795, second row: templates protected using block remapping 32×32 - comparison score 0.1179, third row: original input templates pre-aligned - comparison score 0.1832, fourth row: templates protected using block remapping 32×32 after pre-alignment - comparison score 0.1785.

the UTFVP dataset, it is evident that pre-aligning the finger vein templates significantly improves the results for all the considered BTP approaches, allowing to obtain recognition performances close to those of the baseline unprotected systems. Variances in the EER for block remapping and warping are caused by the actual block selection [32]: if the particular key selects more central blocks, the EER is decreased, while if the particular key selects more border blocks, the EER increases. The highest improvement can be achieved for block remapping and block sizes of 16×16 as well

Table 4
Recognition performance results in terms of EER and 95% confidence intervals for cancelable biometrics schemes applied on the UTFVP database with feature pre-alignment, using 10 different transformation keys for each template.

Scheme	Parameters	GF	IUWT	MC	PC	RLT	WLD
Block Remapping	16x16	4.26%±0.39	4.63%±0.41	**1.62%±0.26**	2.07%±0.29	1.81%±0.25	2.08%±0.31
	32x32	1.76%±0.39	2.22%±0.41	**0.69%±0.25**	1.34%±0.29	0.83%±0.25	1.3%±0.3
	48x48	4.77%±0.42	5.6%±0.43	**1.81%±0.28**	2.64%±0.33	1.99%±0.3	2.42%±0.32
	64x64	1.99%±0.33	2.17%±0.34	**0.7%±0.21**	1.02%±0.24	1.47%±0.26	1.2%±0.26
Block Warping	16 - 6	0.74%±0.18	0.74%±0.18	**0.23%±0.11**	0.65%±0.16	0.74%±0.18	0.65%±0.15
	32 - 12	0.69%±0.18	0.74%±0.18	**0.27%±0.11**	0.56%±0.16	0.83%±0.18	0.37%±0.14
	48 - 18	0.83%±0.19	0.88%±0.19	**0.32%±0.12**	0.6%±0.16	0.75%±0.18	0.55%±0.14
	64 - 24	0.74%±0.18	0.6%±0.18	**0.23%±0.12**	0.41%±0.16	0.6%±0.18	0.32%±0.14
Bloom Filters	48/10/7/14	7.04%±0.75	2.39%±0.31	2.23%±0.33	**1.25%±0.26**	2.22%±0.32	3.62%±0.4

Table 5
Recognition performance results in terms of EER and 95% confidence intervals for cancelable biometrics schemes applied on the SDUMLA-HMT database with feature pre-alignment, using 10 different transformation keys for each template.

Scheme	Parameters	GF	IUWT	MC	PC	RLT	WLD
Block Remapping	16x16	30.8%±0.48	13.5%±0.4	**11.7%±0.39**	17.0%±0.42	13.9%±0.39	10.9%±0.39
	32x32	33.2%±0.47	19.1%±0.43	**16.9%±0.41**	20.0%±0.43	18.5%±0.42	18.5%±0.43
	48x48	28.5%±0.46	11.4%±0.39	**8.04%±0.39**	10.9%±0.4	13.6%±0.39	9.54%±0.39
	64x64	14.4%±0.41	3.73%±0.23	**2.16%±0.21**	3.66%±0.23	9.23%±0.32	3.82%±0.24
Block Warping	16 - 6	12.72%±0.33	3.98%±0.21	**3.05%±0.18**	3.94%±0.21	9.25%±0.3	4.34%±0.21
	32 - 12	11.2%±0.35	4.17%±0.21	**2.86%±0.18**	4.15%±0.22	10.91%±0.32	4.28%±0.23
	48 - 18	12.46%±0.34%	5.51%±0.27%	**3.5%±0.22%**	5.65%±0.27%	12.59%±0.34%	5.63%±0.27%
	64 - 24	9.5%±0.33	4.05%±0.2	**2.56%±0.16**	4.07%±0.2	10.94%±0.32	4.66%±0.22
Bloom Filters	22/4/7/16	19.8%±0.442	10.7%±0.31	11.7%±0.32	12.9%±0.34	24.35%±0.44	**6.1%±0.24**

as 32×32 and for the Bloom filters using PC features. With pre-alignment, applying Bloom filters to PC features outperforms almost all block-remapping combinations, with the use of MC features being the only exception. For block warping, a consistent recognition accuracy improvement is achieved for all feature types and all tested transformation parameters. On the SDUMLA dataset, the same general trend as on the UTFVP one holds as well. Pre-alignment considerably improves the results for block remapping and block warping for all combinations. For the Bloom filter approach there is an improvement as well, but it is less significant than for block remapping and block warping. The highest improvement can again be achieved for the block remapping approach using a block size of 16×16. The best performance for the Bloom filters is again achieved using WLD features. The best overall performance is achieved using block warping with a block size of 64 and a maximum offset of 24 for the MC features.

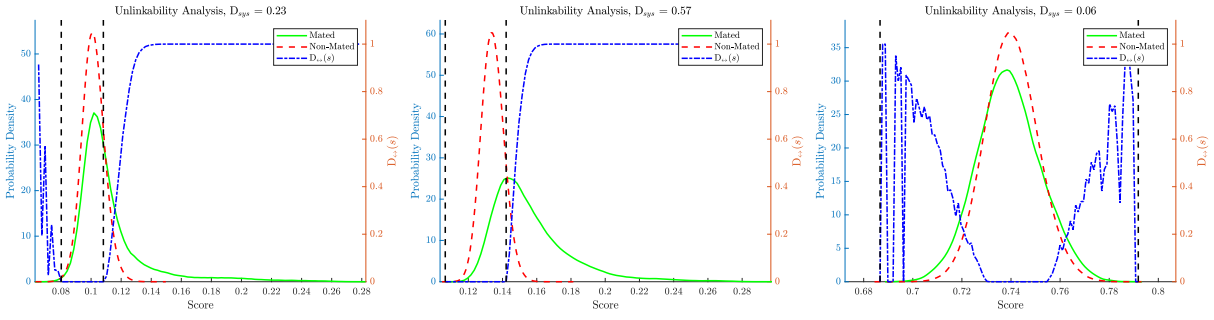

(a) (b) (c)

Figure 9: Unlinkability analysis on the UTFVP dataset. Mated-sample (solid green) and non-mated-sample (dashed red) score distibutions for protected templates. The blue curve represents the score-wise linkability measure $D_{\leftrightarrow}(s)$, and $D_{\leftrightarrow}^{sys}$ gives an estimation of the overall system linkability. (a): Block remapping using MC and $B = 64$; (b): Block warping using MC, $G = 64$ and $O = 24$; (c): Bloom filters.

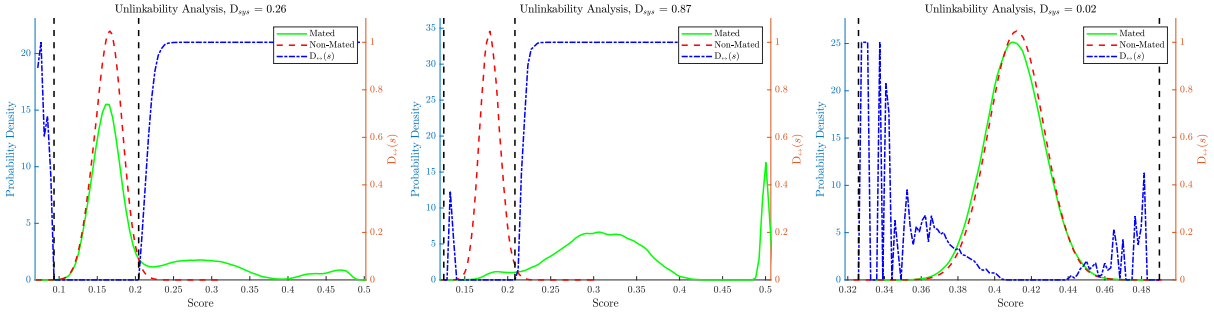

(a) (b) (c)

Figure 10: Unlinkability analysis on the SDUMLA-HMT dataset. Mated-sample (solid green) and non-mated-sample (dashed red) score distibutions for protected templates. The blue curve represents the score-wise linkability measure $D_{\leftrightarrow}(s)$, and $D_{\leftrightarrow}^{sys}$ gives an estimation of the overall system linkability. (a): Block remapping using MC and $B = 64$; (b): Block warping using MC, $G = 64$ and $O = 24$; (c): Bloom filters.

5.3. Security Analysis

In this section we evaluate the compliance of the proposed systems with the requirements established in the ISO/IEC 24745 standard on biometric information protection [17]. Specifically, in Section 5.3.1 we thoroughly analyse the unlinkability of the employed template protection schemes, while their irreversibility is experimentally evaluated in Section 5.3.2.

5.3.1. Unlinkability Analysis

The unlinkability analysis is performed for all three applied cancelable biometrics schemes, on both datasets and for all six feature types. The local unlinkability measure $D_{\leftrightarrow}(s)$ is computed for selected combinations, with the obtained results depicted as blue curves in Figure 9 for UTFVP, and in Figure 10 for the SDUMLA-HMT dataset. The global measures $D_{\leftrightarrow}^{sys}$ obtained for all the considered parameter combinations are instead listed in Table 6 and Table 7 for the UTFVP and the SDUMLA-HMT databases, respectively. As it may be observed, especially in the

Table 6
Unlinkability results in terms of $D_{\leftrightarrow}^{sys}$ for all cancelable schemes and all six feature types on the UTFVP dataset.

Scheme	GF	IUWT	MC	PC	RLT	WLD
BM 16	0.007	0.014	0.064	0.018	0.042	0.081
BM 32	0.016	0.02	0.079	0.09	0.033	0.076
BM 48	0.139	0.129	0.198	0.115	0.083	0.235
BM 64	0.112	0.116	0.229	0.078	0.11	0.176
BW 16/6	0.829	0.891	0.918	0.886	0.857	0.807
BW 32/12	0.435	0.498	0.658	0.572	0.52	0.501
BW 48/18	0.421	0.451	0.562	0.509	0.44	0.468
BW 64/24	0.412	0.441	0.566	0.502	0.48	0.468
Bloom Filter	0.344	0.029	0.063	0.024	0.027	0.031

Table 7
Unlinkability results in terms $D_{\leftrightarrow}^{sys}$ for all cancelable schemes and all six feature types on SDUMLA dataset.

$D_{\leftrightarrow}^{sys}$	GF	IUWT	MC	PC	RLT	WLD
BM 16	0.013	0.034	0.054	0.028	0.018	0.038
BM 32	0.047	0.083	0.158	0.104	0.056	0.123
BM 48	0.035	0.026	0.027	0.02	0.024	0.023
BM 64	0.092	0.242	0.256	0.238	0.193	0.245
BW 16/6	0.736	0.506	0.445	0.503	0.397	0.442
BW 32/12	0.462	0.403	0.395	0.427	0.289	0.374
BW 48/18	0.347	0.462	0.527	0.549	0.414	0.518
BW 64/24	0.359	0.847	0.865	0.854	0.764	0.843
Bloom Filter	0.031	0.043	0.019	0.022	0.046	0.043

case of block warping and MC features, there is no big overlap between the mated (green) and non-mated (red) score distributions (area under the curve within the overlapping parts is considerably smaller than outside the overlapping parts), being $p\left(H_m|s\right) > p\left(H_{nm}|s\right)$ for $s > 0.15$ in Figure 9b for UTFVP and Figure 10b for SDUMLA, and in case of block remapping for $s > 0.22$ in Figure 9a for the UTFVP and Figure 10a.

No overlap corresponds to separable scores, i.e., linkable templates. Hence, for those intervals with no overlap $D_{\leftrightarrow}(s) = 1$, the templates are fully linkable. For block warping, since most the mated instances score distributions lies in the aforementioned score interval, the global linkability of the systems $D_{\leftrightarrow}^{sys}$ is 0.57, thereby showing that the considered scheme fulfills the unlinkability requirement only partially. The block remapping scheme instead shows proper unlinkability for smaller block sizes, with a notable unlinkability degradation for larger block values. The lowest linkability measures are obtained when employing Bloom filters as BTP scheme. The unlinkability requirement

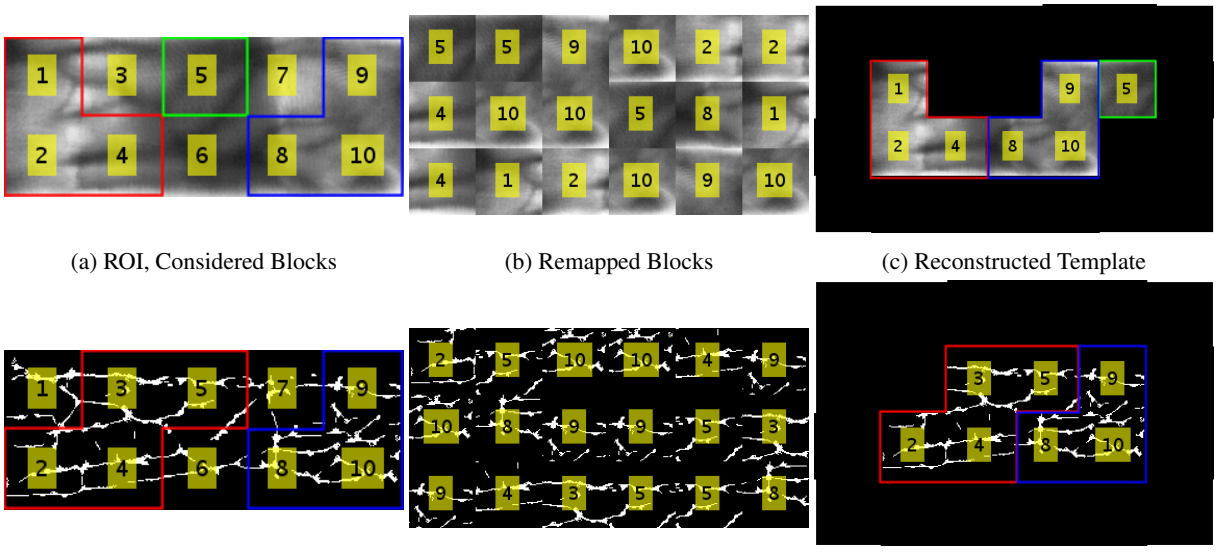

(a) ROI, Considered Blocks (b) Remapped Blocks (c) Reconstructed Template

(d) MC: ROI, Considered Blocks (e) MC: Remapped Blocks (f) MC: Reconstructed Template

Figure 11: Template reconstruction using PuzzleMultisolver [31], for the original image (top) and the corresponding MC features (bottom). In each row: ROI, labeled blocks and used blocks

is therefore satisfied only by the block-remapping and the Bloom-filter schemes, whereas the block-warping approach is not suitable for template protection in terms of unlinkability.

5.3.2. Irreversibilty Analysis

As already discussed in the previous section, the block-remapping approach shows proper unlinkability performance to be considered as BTP scheme. Its irreversibility is therefore here also analysed in detail, exploiting the automated square jigsaw puzzle solver algorithm introduced in Section 4.2. The experiment uses the code provided with the original article [31][4]. Figure 11 depicts the process. The left column of Figure 11 reports the original template. The blocks considered during remapping are grouped into regions of connected blocks, and their outline is highlighted. The middle column shows the block remapped images. This image consists only of the considered blocks. The right column shows the square jigsaw puzzle solver reconstruction results. Again, this image consists only of the considered blocks. In a successful reconstruction, all block regions are restored. Due to the omitted blocks, an exact arrangement of the regions is not always possible. The blocks are marked with the same numbers across all three images. The amount of information from the original template which can be restored is hereby directly linked to the irreversibility property of the template protection scheme, with the highest possible amount of reconstructed data r being $r = \frac{B_{pt}}{B_{ot}}$, where B_{pt} is the number of blocks considered in the protected template, and B_{ot} is the number of blocks contained in the original template.

The irreversibility analysis is performed for block remapping on both the UTFVP and the SDUMLA datasets.

[4] https://github.com/ZaydH/sjsu_thesis

Table 8
Irreversibilty Analysis for Block Remapping on the UTFVP dataset.

Reconstructed Pairs						
Block size	GF	IUWT	MC	PC	RLT	WLD
16x16	17.8%	16.6%	9.24%	16.8%	9.61%	9.57%
32x32	72.4%	71.5%	47.3%	69.3%	47.5%	42.0%
48x48	90.0%	89.7%	76.6%	90.5%	77.0%	71.1%
64x64	93.1%	94.0%	87.9%	94.3%	90.4%	84.7%
Max. Reconstructed Region Size						
Block size	GF	IUWT	MC	PC	RLT	WLD
16x16	4.31%	4.18%	2.94%	4.27%	2.90%	3.02%
32x32	54.9%	53.7%	30.6%	51.4%	32.2%	28.1%
48x48	90.3%	89.7%	73.4%	89.5%	74.0%	65.5%
64x64	93.4%	95.1%	88.1%	95.7%	93.2%	84.7%
Perfect Reconstruction						
Block size	GF	IUWT	MC	PC	RLT	WLD
16x16	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
32x32	0.24%	0.17%	0.00%	0.11%	0.00%	0.00%
48x48	36.8%	35.9%	11.1%	38.3%	11.7%	4.82%
64x64	64.6%	69.2%	46.1%	70.6%	55.9%	39.7%

Different block sizes (16×16, 32×32, 48×48 and 64×64) and all six feature types are evaluated. Table 8 and Table 9 list the results of the puzzle solver approach, averaged over each single run (key), and then again over all the 10 different keys, for the UTFVP and the SDUMLA-HMT databases, respectively. The reported values are relative to the maximum possible amount r of data that can be reconstructed, that is, if only 7 out of 10 blocks are considered, and the value in the table is 100%, this means that $100\% \cdot \frac{7}{10} = 70\%$ of the total unprotected template has been successfully reconstructed. The results of both datasets show that the reconstruction performance increases with the size of the blocks used. For UTVFP with a block size of 16×16 pixels only $< 20\%$ of the possible pairs are correctly identified, and the maximum possible region is reconstructed only for less than 5% of the templates. For 64×64 blocks, 85-95% of all available pairs are reconstructed. The largest possible region is correctly detected for 86-96% of all templates. When comparing the number of perfect reconstructions, the difference is even bigger. While no perfect reconstruction is obtained for 16×16 blocks, this achievement is accomplished for $> 40\%$ of the 64×64 templates. In the best case, when using PC features, perfect reconstructions are achieved for even 70% of all templates. It is not surprising that the jigsaw reconstruction performs better for larger blocks, since blocks with longer borders provide more information for the reconstruction.

Table 9
Irreversibilty Analysis for Block Remapping on the SDUMLA dataset.

Reconstructed Pairs						
Block size	GF	IUWT	MC	PC	RLT	WLD
16x16	16.7%	39.8%	30.6%	39.4%	31.1%	30.6%
32x32	71.7%	90.3%	83.2%	87.6%	83.7%	82.5%
48x48	89.4%	98.2%	95.3%	98.7%	95.6%	94.9%
64x64	94.3%	100.0%	94.0%	73.0%	100.0%	98.9%
Max. Reconstructed Region Size						
Block size	GF	IUWT	MC	PC	RLT	WLD
16x16	4.40%	27.3%	20.4%	26.0%	22.3%	20.8%
32x32	55.9%	93.3%	86.0%	89.4%	87.3%	85.3%
48x48	89.1%	98.2%	95.6%	98.7%	95.9%	95.3%
64x64	97.1%	100.0%	97.0%	86.5%	100.0%	99.5%
Perfect Reconstruction						
Block size	GF	IUWT	MC	PC	RLT	WLD
16x16	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
32x32	0.61%	59.2%	44.0%	51.5%	38.1%	39.9%
48x48	37.2%	92.8%	85.8%	95.1%	85.0%	84.2%
64x64	74.0%	100.0%	94.0%	73.0%	100.0%	98.9%

Table 10
Number of blocks used for the different block remapping settings on the UTFVP and SDUMLA dataset.

Dataset	Block size	GF	IUWT	MC	PC	RLT	WLD
UTFVP	16x16	168 (8x21)	168 (8x21)	168 (8x21)	168 (8x21)	168 (8x21)	168 (8x21)
	32x32	44 (4x11)	44 (4x11)	44 (4x11)	44 (4x11)	44 (4x11)	44 (4x11)
	48x48	21 (3x7)	21 (3x7)	21 (3x7)	21 (3x7)	21 (3x7)	21 (3x7)
	64x64	12 (2x6)	12 (2x6)	12 (2x6)	12 (2x6)	12 (2x6)	12 (2x6)
SDUMLA	16x16	160 (8x20)	40 (4x10)	40 (4x10)	40 (4x10)	40 (4x10)	40 (4x10)
	32x32	40 (4x10)	10 (2x5)	10 (2x5)	10 (2x5)	10 (2x5)	10 (2x5)
	48x48	21 (3x7)	8 (2x4)	8 (2x4)	8 (2x4)	8 (2x4)	8 (2x4)
	64x64	10 (2x5)	3 (1x3)	3 (1x3)	3 (1x3)	3 (1x3)	3 (1x3)

Another important factor for the achieved reconstruction rates is the number of available blocks. The results on the SDUMLA show that less available blocks increase the reconstruction rates. With the exception of GF, where the size of the templates is similar to those of UTFVP, the templates extracted from SDUMLA have roughly half the size of their UTFVP counterparts (with respect to their dimensions). Table 10 depicts the number of available blocks for the respective configurations. The experimental results show that a reduced number of blocks leads to better

reconstruction rates. While on the SDUMLA for GF the rates are similar to those on UTFVP, for the other feature types (IUWT, MC, PC, RLT and WLD) the success rates are noticeable better on SDUMLA than on the UTFVP. In fact, when considering SDUMLA, 30-40% of all available pairs are correctly detected for 16×16 blocks, while for UTFVP the same holds true for only less than 20%. The largest region is retrieved for 20-30% of SDMULA templates, while for UTFVP the jigsaw solver is able to do so only for less than 5%. The same behavior can be seen for the other block sizes. For 64×64 blocks on SDUMLA, two feature types (IUWT and RLT) even achieve a perfect reconstruction (the jigsaw algorithm is capable of placing all pairs correctly in the reconstruction) for all available templates. The only configuration that does not behave as described above is PC features, for a block size of 64×64 pixels, on SDUMLA. For this setting, despite the lower number of available blocks, the reconstruction results are worse than on the UTFVP dataset.

When comparing the reconstruction results with the recognition performance results in Section 5.2 it turns out, that for settings with better recognition rates, also the reconstruction rates of the jigsaw puzzle solver are higher and therefore lower irreversibility. The achieved reconstruction rates indicate that once an attacker gets hold of the original template and the protected one, it is possible to reconstruct the key , i.e., the mapping information, which in turn poses another threat for this kind of template protection scheme in case a system-dependent key is used. Keeping in mind that the utilised puzzle solver is not optimised for the reconstruction of finger vein templates, even better reconstruction rates are expected to be achieved with an optimised version of the puzzle solver.

For BTP schemes based on Bloom filters, the success probability for an attack trying to recover the original unprotected features from their protected representation can be estimated as $\overline{nSeq}^{nBlocks}$, being $\overline{nSeq}$ the average number of possible sequences resulting in a single Bloom filter, defined in eq. (5). In the considered tests, the success probabilities for guessing the original unprotected templates range from 10^{-192} to 10^{-23}, therefore confirming the irreversibility of the templates.

5.4. Results Discussion and Summary

In terms of recognition capabilities, the block warping scheme performs best, as it achieves the highest recognition performance in terms of EER, followed by the block remapping scheme, with the Bloom filters leading to the worst performance. The different feature types have an impact on the recognition performance and the security as well, but the general trend remains the same among all six tested feature types.

For both block remapping and warping there is a general trend of recognition performance improving with increasing block sizes. Actually, there is a trade-off between recognition performance and security, in terms of unlinkability and irreversibility, observed for both block remapping and block warping approaches: changing the employed transformation parameters, the higher the recognition performance, the lower the level of security, and vice versa.

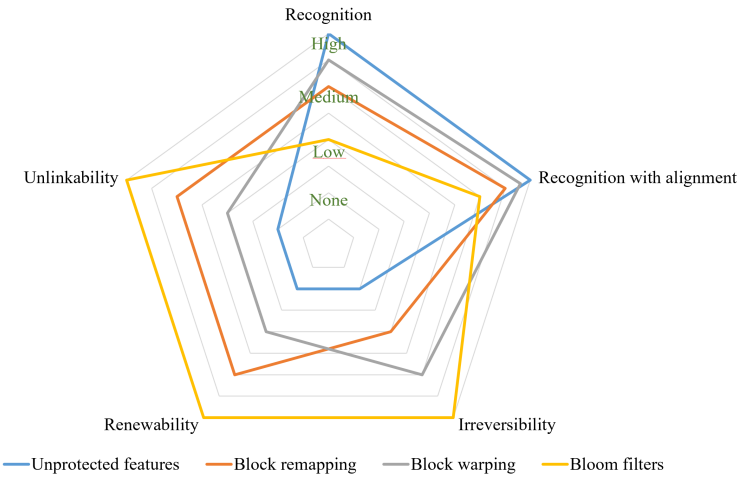

Figure 12: Qualitative comparison of the considered cancelable biometrics approaches, taking into account recognition performance, irreversibility, unlinkability, and renewability.

Regarding the Bloom-filter-based approach, with respect to its application to spectral minutiae representations presented in [10], our results for the protected templates without performing the pre-alignment are less performing. However, it is worth observing that the detection and extraction of reliable minutiae points for finger veins is a difficult and error-prone task. Hence, binary representations are more commonly used, as they guarantee more reliable and stable results (our baseline EER of 0.36% on the UTFVP dataset compared to the baseline EER of 1.5% as reported in [10]). On the other hand, the employed binary representations have the aforementioned inherent alignment problems, as for all block-based template protection schemes. This is confirmed by the fact that our results, in case the pre-alignment strategy is employed, are superior to the ones reported in [10] (0.23% EER for Block Warping 64 - 24, and 1.25% EER for Bloom filters, instead of 2.1% EER reported in [10]).

In terms of unlinkability, block warping achieves the lowest security, thereby unveiling the method inadequacy as BTP scheme. The employed puzzle-solver attack also shows that the block remapping scheme is not secure enough, since its irreversibility solely relies on the amount of blocks which are not considered, and even the key can be reconstructed under certain circumstances. Hence, in terms of security, the Bloom filter approach remains the only effective solution.

A graphical summary of the strengths and weaknesses of the considered cancelable biometrics approaches is provided in Figure 12, where the employed methods are qualitatively compared in terms of achievable recognition performance, irreversibility, unlinkability, and renewability. As it can be seen, block remapping and block warping provide recognition rates comparable with those of unprotected approaches, yet only a limited irreversibility and unlinkability can be obtained. Conversely, using Bloom filters as template protection scheme allows attaining the best results in terms of irreversibility and unlinkability, at the cost of a higher recognition performance degradation, especially in case it is not possible to properly align the original features to be compared. Consequently, the decision

regarding which kind of template protection scheme should be used depends on the specific requirements: if recognition performance is more important, block remapping or warping should be applied, whereas if security is the main concern, the Bloom filter approach is a better choice.

The effects of template misalignments on the effectiveness of Bloom filters are shown in Figure 13, which reports the average mated scores computed for increasing amounts of translations and rotations between the compared templates. The reported results show that, while template translations along the x-axis can be fairly tolerated, differences in alignments along the y-axis significantly affect the recognition rates. Limited rotation displacements remarkably worsen the similarity scores computed from mated comparisons. A pre-alignment approach is therefore highly relevant for implementing cancelable biometrics approaches usable in practical applications. The results in Figure 13 show that rotational displacements are specifically relevant, and should be therefore carefully addressed. The solution here proposed is actually beneficial for all the block-based cancelable biometrics schemes, being able to considerably improve the performance. Yet, since such proof-of-concept alignment method does not comply with BTP principles, there is need for further investigating a universal alignment method, based on some kind of finger landmarks.

6. Conclusion

In this article we evaluated three different cancelable biometrics schemes for finger vein recognition, namely block remapping, block warping, and Bloom filters. Six different feature extractors of well-established vein recognition schemes, producing binary templates, were utilised to generate the unprotected templates. These templates were then protected using the aforementioned cancelable biometrics schemes. In addition, a pre-alignment approach prior to the application of the cancelable schemes is proposed and tested. The evaluation was conducted on two well-known finger vein datasets, the UTFVP and the SDUMLA-HMT databases. Recognition performance, unlinkability, and irreversibility were evaluated.

Block remapping and block warping, in combination with the pre-alignment, achieved the best results in terms of recognition performance. However, block remapping is not secure enough, as it turned out that its unlinkability, as well as its irreversibility, are rather low (the irreversibility solely depends on the number of disregarded blocks). Block warping has a low unlinkability as well. Hence, only the Bloom filter approach is suitable in terms of security, being able to withstand RMAs too. In combination with the pre-alignment, it achieves an acceptable recognition performance, although this recognition performance is still much worse than the baseline one. However, an application in a multi-modal biometric system might be feasible (higher error rates can be compensated by the other deployed modalities).

In general, poor recognition performance is achieve without pre-alignment. Thus, an accurate, universal pre-alignment, which does not require the unprotected templates to be present in the system, is necessary in order to em-

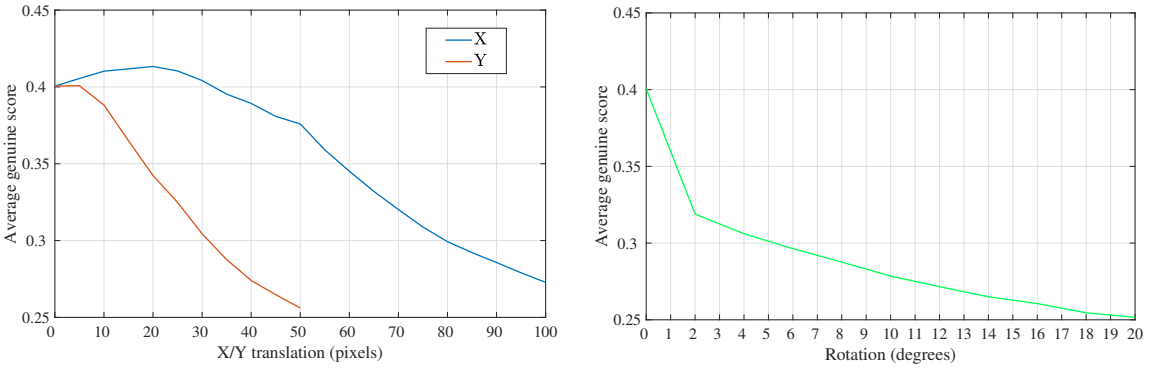

Figure 13: Effects of template misalignments on the recognition performance achievable using Bloom filters, evaluated on the UTFVP database: (left) Translations, (right) Rotations. Higher score values correspond to templates being more similar.

ploy a well-performing (in terms of recognition accuracy) template protection scheme. Such universal pre-alignment must not be based on the extracted templates but only on some "landmarks" present in the input images, and shall be able to align one input sample to a "universal" position without needing a reference sample/template.

In our future work we will aim for further performance improvements of the Bloom filter approach, as this method turned out to be the most beneficial one in terms of security and privacy. The main limitations of the recognition performance for the Bloom filter approach are the shifts/rotations present in the input samples. Hence, the first step to improve the performance is to come up with the aforementioned universal pre-alignment. The next step is to fine-tune and optimise the parameters of the template protection/feature transformation, in order to achieve the best possible recognition performance while still maintaining an adequate level of template protection. The last step regards run-time performance improvement, especially in the comparison step of the Bloom-filter approach. Furthermore, we will evaluate other variants of block remapping and block warping, like remapping including shifts in the blocks and recursive remapping, as well as other strategies to derive the warped grid in the warping approach. The main limitation of all those block-based approaches is the trade-off between security and recognition performance. The lower the block size, the higher the number of blocks and the higher the security, but the lower the recognition performance, again mostly due to shifts/rotations in the input data. Hence, also here a universal pre-alignment approach would be needed. Eventually, it would be highly desirable to find the sweet spot between recognition performance and security for the different block remapping and warping variants.

Acknowledgements

This work has received funding from the European Unions Horizon 2020 research and innovation program under grant agreement No. 700259. It has received further funding by the Austrian Science Fund FWF and funding by the

Salzburg state government, project no. P32201 - Advanced Methods and Applications for Fingervein Recognition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

CRediT authorship contribution statement

Christof Kauba 1234567891

1234567891Please cite this work as: Christof Kauba, Emanuela Piciucco, Emanuele Maiorana, Marta Gomez-Barrero, Bernhard Prommegger, Patrizio Campisi and Andreas Uhl, Towards Practical Cancelable Biometrics for Finger Vein Recognition, Elsevier Information Sciences, in press, 2021: Conceptualization; Data curation; Investigation; Methodology; Software; Validation; Visualization; Roles/Writing - original draft; Writing - review & editing. **Emanuela Piciucco:** Conceptualization; Investigation; Methodology; Software; Roles/Writing - original draft; Writing - review & editing.. **Emanuele Maiorana:** Methodology; Supervision; Validation; Roles/Writing - original draft; Writing - review & editing.. **Marta Gomez-Barrero:** Investigation; Methodology; Software; Roles/Writing - original draft; Writing - review & editing.. **Bernhard Prommegger:** Data curation; Methodology; Software; Visualization; Roles/Writing - original draft. **Patrizio Campisi:** Conceptualization; Formal analysis; Investigation; Methodology; Roles/Writing - original draft; Writing - review & editing.. **Andreas Uhl:** Conceptualization; Funding acquisition; Project administration; Resources; Supervision; Writing - review & editing.

References

[1] Abe, N., Yamada, S., Shinzaki, T., 2015. Irreversible fingerprint template using minutiae relation code with Bloom filter, in: Proc. Int. Conf. on Biometrics Theory, Applications and Systems (BTAS), pp. 1–7.

[2] Bianchi, T., Piva, A., 2012. Image forgery localization via block-grained analysis of jpeg artifacts. IEEE Transactions on Information Forensics and Security 7, 1003–1017.

[3] Boult, T.E., Scheirer, W.J., Woodworth, R., 2007. Revocable fingerprint biotokens: accuracy and security analysis, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–8.

[4] Cho, T.S., Avidan, S., Freeman, W.T., 2010. A probabilistic image jigsaw puzzle solver, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'10), IEEE. pp. 183–190.

[5] Choi, J.H., Song, W., Kim, T., Lee, S.R., Kim, H.C., 2009. Finger vein extraction using gradient normalization and principal curvature. IS&T/SPIE Electronic Imaging 7251, 1–9.

[6] Debiasi, L., Kirchgasser, S., Prommegger, B., Uhl, A., Grudzień, A., Kowalski, M., 2019. Biometric template protection in the image domain using non-invertible grey-scale transforms, in: 2019 IEEE International Workshop on Information Forensics and Security (WIFS), IEEE. pp. 1–6.

[7] Galbally, J., Ross, A., Gomez-Barrero, M., Fierrez, J., Ortega-Garcia, J., 2013. Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms. Computer Vision and Image Understanding 117, 1512–1525.

[8] Gomez-Barrero, M., Galbally, J., Rathgeb, C., Busch, C., 2018a. General framework to evaluate unlinkability in biometric template protection systems. IEEE Transactions on Information Forensics and Security 13, 1406–1420.

[9] Gomez-Barrero, M., Rathgeb, C., Galbally, J., Busch, C., Fierrez, J., 2016. Unlinkable and irreversible biometric template protection based on Bloom filters. Information Sciences 370-371, 18–32.

[10] Gomez-Barrero, M., Rathgeb, C., Li, G., Ramachandra, R., Galbally, J., Busch, C., 2018b. Multi-biometric template protection based on Bloom filters. Information Fusion 42, 37–50.

[11] Guo, H., Burrus, C.S., 1996. Convolution using the undecimated discrete wavelet transform, in: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICCASP'96), IEEE. pp. 1291–1294.

[12] Hämmerle-Uhl, J., Pschernig, E., Uhl, A., 2009. Cancelable iris biometrics using block re-mapping and image warping, in: Information Security. Springer, pp. 135–142.

[13] Hermans, J., Mennink, B., Peeters, R., 2014. When a Bloom filter is a doom filter: Security assessment of a novel iris biometric template protection system, in: IEEE International Conference of the Biometrics Special Interest Group (BIOSIG).

[14] Hine, G.E., Maiorana, E., Campisi, P., 2017. A zero-leakage fuzzy embedder from the theoretical formulation to real data. IEEE Transactions on Information Forensics and Security 12.

[15] Hirata, S., Takahashi, K., 2009. Cancelable biometrics with perfect secrecy for correlation-based matching, in: Proceedings of the International Conference on Biometrics (ICB), IEEE. pp. 868–878.

[16] Huang, B., Dai, Y., Li, R., Tang, D., Li, W., 2010. Finger-vein authentication based on wide line detector and pattern normalization, in: Proceedings of the 20th International Conference on Pattern Recognition (ICPR'10), IEEE. pp. 1269–1272.

[17] ISO/IEC JTC1 SC27 Security Techniques, 2011. ISO/IEC 24745:2011. Information Technology - Security Techniques - Biometric Information Protection. International Organization for Standardization.

[18] Jain, A.K., Nandakumar, K., Nagar, A., 2008. Biometric template security. EURASIP Journal on Advances in Signal Processing 2008, 117.

[19] Jenisch, S., Uhl, A., 2011. Security analysis of a cancelable iris recognition system based on block remapping, in: Proceedings of the IEEE International Conference on Image Processing, (ICIP'11), Brussels, Belgium. pp. 3274–3277.

[20] Kelkboom, E.J.C., Breebaart, J., Kevenaar, T., Buhan, I., Veldhuis, R., 2011. Preventing the decodability attack based crossmatching in a fuzzy commitment scheme. IEEE Transactions on Information Forensics and Security 6, 107121.

[21] Kirchgasser, S., Kauba, C., Lai, Y.L., Zhe, J., Uhl, A., 2020. Finger vein template protection based on alignment-robust feature description and index-of-maximum hashing. IEEE Transactions on Biometrics, Behavior, and Identity Science 2, 337–349. doi:10.1109/TBIOM.2020.2981673.

[22] Kirchgasser, S., Kauba, C., Uhl, A., 2019. Cancellable biometrics for finger vein recognition - application in the feature domain, in: Uhl, A., Busch, C., Marcel, S., Veldhuis, R. (Eds.), Handbook of Vascular Biometrics. Springer Nature Switzerland, Cham, Switzerland. chapter 16, pp. 481–506.

[23] Kumar, A., Zhou, Y., 2012. Human identification using finger images. IEEE Transactions on Image Processing 21.

[24] Lee, E.C., Lee, H.C., Park, K.R., 2009. Finger vein recognition using minutia-based alignment and local binary pattern-based feature extraction. International Journal of Imaging Systems and Technology 19, 179–186.

[25] Leng, L., Zhang, J., 2013. Palmhash code vs. palmphasor code. Neurocomputing 108, 1–12.

[26] Li, C., Hu, J., 2014. Attacks via record multiplicity on cancelable biometrics templates. Concurrency and Computation: Practice and

Experience 26, 1593–1605.

[27] Maiorana, E., Campisi, P., Fierrez, J., Ortega-Garcia, J., Neri, A., 2010. Cancelable templates for sequence-based biometrics with application to on-line signature recognition. IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans 40.

[28] Miura, N., Nagasaka, A., Miyatake, T., 2004. Feature extraction of finger-vein patterns based on repeated line tracking and its application to personal identification. Machine Vision and Applications 15.

[29] Miura, N., Nagasaka, A., Miyatake, T., 2007. Extraction of finger-vein patterns using maximum curvature points in image profiles. IEICE transactions on information and systems 90.

[30] Ouda, O., 2021. On the practicality of local ranking-based cancelable iris recognition. IEEE Access 9.

[31] Paikin, G., Tal, A., 2015. Solving multiple square jigsaw puzzles with missing pieces, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'15), pp. 4832–4839.

[32] Piciucco, E., Maiorana, E., Kauba, C., Uhl, A., Campisi, P., 2016. Cancelable biometrics for finger vein recognition, in: Proceedings of the First international workshop on sensing, processing and learning for intelligent machines (SPLINE), pp. 1–5.

[33] Pillai, J.K., Patel, V.M., Chellappa, R., Ratha, N.K., 2011. Secure and robust iris recognition using random projections and sparse representations. Pattern Analysis and Machine Intelligence, IEEE Transactions on 33, 1877–1893.

[34] Quan, F., Fei, S., Anni, C., 2008. Cracking cancelable fingerprint template of Ratha, in: International Symposium on Computer Science and Computational Technology.

[35] Ranbaduge, T., Schnell, R., 2020. Securing Bloom filters for privacy-preserving record linkage, in: ACM International Conference on Information & Knowledge Management.

[36] Ratha, N.K., Connell, J.H., Bolle, R.M., 2001. Enhancing security and privacy in biometrics-based authentication systems. IBM systems Journal 40, 614–634.

[37] Rathgeb, C., Breitinger, F., Busch, C., Baier, H., 2014. On application of Bloom filters to iris biometrics. IET Biometrics 3, 207–218.

[38] Rathgeb, C., Uhl, A., 2011. A survey on biometric cryptosystems and cancelable biometrics. EURASIP Journal on Information Security 2011, 1–25.

[39] Schnell, R., Bachteler, T., Reiher, J., 2009. Privacy-preserving record linkage using Bloom filters. BMC Medical Informatics and Decision Making 9.

[40] Simoens, K., Tuyls, P., Preneel, B., 2009. Privacy weaknesses in biometric sketches, in: Proceedings of the 30th IEEE Symposium on Security and Privacy, IEEE. pp. 188–203.

[41] Ton, B., Veldhuis, R., 2013. A high quality finger vascular pattern dataset collected using a custom designed capturing device, in: Proceedings of the International Conference on Biometrics (ICB'13), IEEE. pp. 1–5.

[42] Uludag, U., Pankanti, S., Prabhakar, S., Jain, A.K., 2004. Biometric cryptosystems: issues and challenges. Proceedings of the IEEE 92, 948–960.

[43] Vidanage, A., Ranbaduge, T., Christen, P., Schnell, R., 2019. Efficient pattern mining based cryptanalysis for privacy-preserving record linkage, in: IEEE International Conference on Data Engineering (ICDE).

[44] Wang, S., Hu, J., 2014. Design of alignment-free cancelable fingerprint templates via curtailed circular convolution. Pattern Recognition 47, 1321–1329.

[45] Wencheng, Y., Song, W., Jiankun, H., Guanglou, Z., Valli, C., 2018. A fingerprint and finger-vein based cancelable multi-biometric system. Pattern Recognition 78, 242–251.

[46] Wolberg, G., 1998. Image morphing: a survey. The visual computer 14, 360–372.

[47] Yin, Y., Liu, L., Sun, X., 2011. SDUMLA-HMT: A multimodal biometric database, in: Sun, Z., Lai, J., Chen, X., Tan, T. (Eds.), Biometric Recognition: 6th Chinese Conference, CCBR 2011, Beijing, China, December 3-4, 2011. Proceedings, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 260–268.

[48] Zhang, J., Yang, J., 2009. Finger-vein image enhancement based on combination of gray-level grouping and circular Gabor filter, in: Proceedings of the International Conference on Information Engineering and Computer Science (ICIECS'09), IEEE. pp. 1–4.

[49] Zhao, J., Tian, H., Xu, W., Li, X., 2009. A new approach to hand vein image enhancement, in: Proceedings of the Second International Conference on Intelligent Computation Technology and Automation (ICICTA'09), IEEE. pp. 499–501.

[50] Zuiderveld, K., 1994. Contrast limited adaptive histogram equalization, in: Heckbert, P.S. (Ed.), Graphics Gems IV. Morgan Kaufmann, pp. 474–485.

Christof Kauba is a post-doc researcher with the Department of Computer Sciences, University of Salzburg, Austria. In 2018, he received his PhD degree in applied information technology from the University of Salzburg where he also pursued his B.Eng. and MSc in 2013 and 2015, respectively. His research interests include image and video processing, image forensics and biometrics, especially biometric sensor design as well as finger- and hand vein biometrics.

Emanuela Piciucco received the bachelor's degree in Electronic Engineering (cum laude) in 2013, and the master's degree in Information and Communication Technology Engineering (cum laude) in 2016, at Roma Tre University, Rome, Italy, where she received her PhD in Applied Electronics in 2020. She is currently a Postdoctoral Researcher with the Section of Applied Electronics, Department of Engineering at Roma Tre University. She was a Visiting Researcher at University of Salzburg, Austria, in 2015, in the framework of the European project ICT COST Action IC1206, and at Telefonica I+D, Barcelona, Spain, in 2017 and 2018, in the framework of the European project ENCASE. Her current research areas are biometric recognition, mainly focusing on vein pattern and EEG biometric identifiers, and physiological signal processing.

Emanuele Maiorana received the Ph.D. degree in biomedical, electromagnetism, and telecommunication engineering with European Doctorate Label from Roma Tre University, Rome, Italy, in 2009. He is currently Assistant Professor with the Section of Applied Electronics, Department of Engineering, Roma Tre University, Rome, Italy. His research interests are in the area of digital signal and image processing, with specific emphasis on biometric recognition. He is an Associate Editor of the IEEE Transactions on Information Forensics and Security. He is the recipient of the Lockheed Martin Best Paper Award for the Poster Track at the IEEE Biometric Symposium 2007, and the Honeywell Student Best Paper Award at the IEEE Biometrics: Theory, Applications and Systems conference 2008.

Bernhard Prommegger received a MSc in Applied Image and Signal Processing in 2014 from a joint degree of the University of Salzburg and University of Applied Sciences Salzburg and a DI (Austrian equivalent to MSc) in Information Technology and Systems Management at the University of Applied Sciences Salzburg in 2015. He is currently pursuing a PhD degree in Applied Image and Signal Processing at the Department of Computer Sciences, University of Salzburg where he is a research assistant. His main research interest is in vascular biometrics, especially multi-perspective finger vein biometrics.

Marta Gomez-Barrero is a Professor for IT-Security and technical data privacy at the Hochschule Ansbach, in Germany. Between 2016 and 2020, she was a postdoctoral researcher at the National Research Center for Applied Cybersecurity (ATHENE) - Hochschule Darmstadt, Germany. Before that, she received her MSc degrees in Computer Science and Mathematics (2011), and her PhD degree in Electrical Engineering (2016), all from Universidad Autonoma de Madrid, Spain. Her current research focuses on security and privacy evaluations of biometric systems, Presentation Attack Detection (PAD) methodologies, and biometric template protection (BTP) schemes. She has co-authored more than 70 publications, chaired special sessions and competitions at international conferences, she is associate editor for the EURASIP Journal on Information Security, and represents the German Institute for Standardization (DIN) in ISO/IEC SC37 JTC1 SC37 on biometrics. She has also received a number of distinctions, including: EAB European Biometric Industry Award 2015, Best Ph.D. Thesis Award by Universidad Autonoma de Madrid 2015/16, Siew-Sngiem Best Paper Award at ICB 2015, Archimedes Award for young researches from Spanish MECD, and Best Poster Award at ICB 2013.

Patrizio Campisi received the Ph.D. degree in electrical engineering from Roma Tre University, Rome, Italy, where he is currently a Full Professor with the Section of Applied Electronics, Department of Engineering. His current research interests are in the area of biometrics and secure multimedia communications. He was the IEEE SPS Director Student Services (2015 - 2017) and the Chair of the IEEE Technical Committee on Information Forensics and Security (2017 - 2018). He is a member of the IEEE Technical Committee on Information Assurance and Intelligent Multimedia-Mobile Communications, System, Man, and Cybernetics Society, and was a member of the IEEE Certified Biometric Program Learning System Committee. He was the General Chair of the 26th European Signal Processing Conference EUSIPCO 2018, Italy, of the 7th IEEE Workshop on Information Forensics and Security (WIFS) 2015, Italy, and of the 12th ACM Workshop on Multimedia and Security 2010, Italy. He is the Editor of the book Security and Privacy in Biometrics (Springer, 2013). He is a Co-Editor of the books Blind Image Deconvolution: Theory and Applications (CRC press, 2007), and High Dynamic Range Video, Concepts, Technologies and Applications (Academic Press, 2016). He was an Associate Editor and a Senior Associate Editor of the IEEE Signal Processing Letters, and an Associate Editor of the IEEE Transactions on Information Forensics and Security. He is currently Editor-in-Chief of the IEEE Transactions on Information Forensics and Security.

Andreas Uhl is a professor at the Department of Computer Sciences (University of Salzburg), where he heads the Multimedia Processing and Security Lab. His research interests include image and video processing and compression, wavelets, media security, medical imaging, biometrics, and number-theoretical numerics.