A. Brömme, C. Busch, A. Dantcheva, K. Raja, C. Rathgeb and A. Uhl (Eds.): BIOSIG 2020,
Lecture Notes in Informatics (LNI), Gesellschaft für Informatik, Bonn 2020 1

Application of affine-based reconstruction to retinal point patterns

Mahshid Sadeghpour, Arathi Arakala, Stephen A. Davis, Kathy J. Horadam[1]

Abstract: Inverse biometrics that exploit the information of biometric references from comparison scores can compromise sensitive personal information of the users in biometric recognition systems. One inverse biometric method that has been very successful in regenerating face images applies an affine transformation to model the face recognition algorithm. This method is general and could apply to templates extracted from other biometric characteristics. This research proposes two formats to apply this method to spatial point patterns extracted from retina images and tests its performance on reconstructing such sparse templates. The results show that the quality of the reconstructed retina point pattern templates is lower than would be accepted by the system as mated.

Keywords: Retina, Biometric Template Protection, MDS-based Reconstruction, Inverse Biometrics, Affine-based Approximation.

1 Introduction

Irreversibility is one of the two critical requirements emphasised by ISO/IEC [IS11] to protect the *sensitive personal data* in biometric references [EC16]. Among reversibility attacks, those that exploit further information of data subjects, inverse biometrics, are classified into four groups based on the level of knowledge required by the attacker [GBG20]. The first group requires knowledge about the template format. The next group, hill-climbing attacks [Ad04, Ga13, Go14], requires access to the scores issued by the system. The third group requires access to the scores generated by the system as well as a set of imposter samples to be presented to the system. The last group are those attacks which require knowledge of the feature extraction method. The method in the third category, which will be reviewed and applied here, was proposed by Mohanty *et al.* [MSK06, MSK12]. Hereinafter we call it the *MSK* method.

The *MSK* algorithm models a biometric comparison algorithm using the scores issued by the system. An attacker needs to have access to a pool of imposter biometric samples (the *break-in set*) to present to the system. This attack is non-iterative. Having access to such a pool, the attacker can perform most of the attack offline without requiring iterative improvements of a presented sample.

The results in [MSK12] are very convincing in regenerating face images, and outperform hill-climbing attacks [Ad04]. So far, this attack has been tested on face recognition systems. The general consensus is that any biometric recognition system that issues the comparison scores is vulnerable toward the *MSK* attack. We would like to check the threat

[1] Mathematical Sciences, School of Science, RMIT University, Melbourne, mahshid.sadeghpour@rmit.edu.au, arathi.arakala@rmit.edu.au, stephen.davis@rmit.edu.au, kathy.horadam@rmit.edu.au

of this attack on different biometric characteristics, those with sparser representations in specific. To the best of our knowledge, it has not yet been tested on vascular biometric characteristics or on point pattern templates. However, the essence of the attack is general enough to be applied to fixed-length templates of other biometric characteristics.

This paper applies this attack to spatial point pattern templates (of varying lengths) extracted from retinal vascular graphs. It has been shown that retina can be accurately represented and compared using sparse templates comprising locations of feature points instead of the whole image [Ar16]. Our intuition was that this type of sparse biometric template might be more resistant to the *MSK* inverse attack.

We propose two formats for inputting the biometric information of the break-in sets as vectors of the same length. By its construction, the *MSK* algorithm should be able to reconstruct the break-in set very well. If the modelling does not do a good job in regenerating the break-in set, it probably would not successfully reconstruct target references. Our experiments show that neither format can reconstruct the break-in set well, nor can they successfully reconstruct the biometric references. The *MSK* algorithm will be reviewed briefly in section 2. The replication of results of Mohanty *et al.* [MSK06] will be presented in section 2.1. Section 3 gives two methods to apply this attack to retinal spatial point pattern templates and tests their performance reconstructing break-in sets and reference databases. Conclusions and future work appear in section 4.

2 Review of the *MSK* algorithm

The *MSK* algorithm approximates any comparison algorithm by an affine transform, given sufficiently many pairwise Euclidean distances between a pool of biometric samples (break-in set) $X = \{X_1, \ldots, X_K\}$. The pairwise distances between the break-in set samples are used to define an affine space where the representations of break-in set samples in the affine space, $\mathbf{Y} = \{\mathbf{Y}_1, \ldots, \mathbf{Y}_K\}$, will have the same pairwise distances as in the Euclidean space. To find the approximating affine transformation $\mathbf{A}$ that maps the break-in set samples $X_i,\ 1 \leq i \leq K$, to the modeled vectors $\mathbf{Y}_i,\ 1 \leq i \leq K$, the *MSK* algorithm inputs the data of break-in templates $\mathbf{X}_i,\ 1 \leq i \leq K$, as columns of a matrix $\mathbf{X}$. Thus, the break-in templates are required to have the same length. For an unknown target sample X_t, if an attacker has access to the distance vector $d' = (d'_1, \ldots, d'_K)$ between X_t and the break-in samples, he can locate the transformed target $\mathbf{Y}_z$ in the affine space. Using a pseudo-inverse $\mathbf{A}^\dagger$ of affine approximation $\mathbf{A}$, the attacker can then reconstruct the target template $\mathbf{X}_z$ from its point $\mathbf{Y}_z$ in the affine space. Full details can be found in [MSK12].

2.1 Reproducing the results of face image reconstruction

We converted the *MSK* Matlab code in [MSK06] to R code, then used it to model the PCA-based face recognition system [TP91] using FERET face database [Ph98, Ph00]. The underlying face recognition system applies Mahalanobis Cosine distance to compare the eigen-faces [Be03]. This experiment is performed to confirm that the R code is capable of

reproducing Mohanty *et al.*'s results in [MSK06]. The break-in set consists of 596 samples of 149 individuals from FERET. The images are 150×130 pixels, resulting in face image vectors of length $19,500$.

First, we tried to reconstruct the break-in set members. From an attacker's point of view, a successful modelling of the biometric algorithm would result in reconstructed break-in samples with high quality. The reason is that when a break-in sample is considered as the target, its distance vector to the break-in samples has 0 in one coordinate, as the target already exists in the break-in set. Figure 1a shows two of the original break-in samples from FERET on the top row and their corresponding reconstructed faces on the bottom row.

We then used the break-in set to reconstruct samples from 100 different target biometric references (BRs) in the FERET dataset and presented the reconstructed face images to the PCA-based system as probe samples.

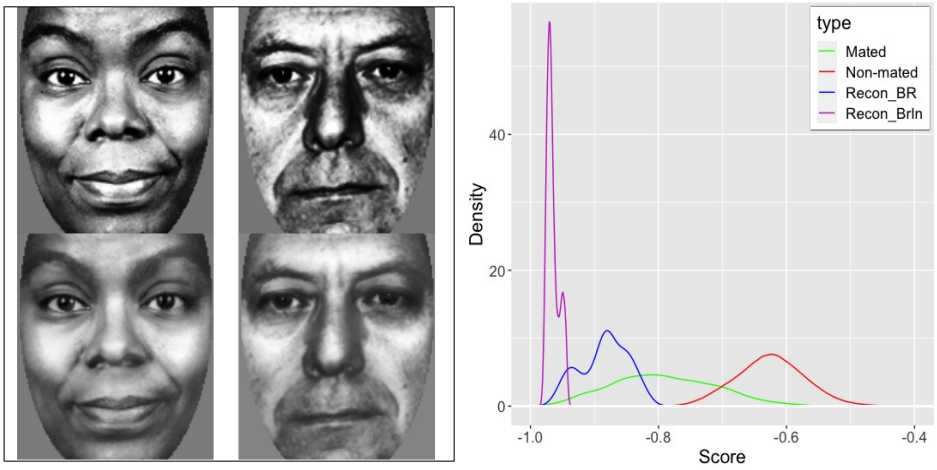

(a) (b)

Fig. 1. a) Instances of some original faces from break-in set (top row) and their corresponding reconstructed faces (bottom row). b) Mahalanobis Cosine score distributions for mated (green), non-mated (red), and *MSK*-reconstructed samples (purple, reconstructed break-in set; and blue, reconstructed target BRs)

Figure 1b illustrates the distributions of Mahalanobis scores for mated, non-mated, and reconstructed samples in the PCA-based face recognition system. The distribution of scores achieved by comparing reconstructed reference samples to the original samples (the blue curve in Figure 1b) shows that all the reconstructed faces could be accepted by the system, which confirms the results in [MSK06]. The scores obtained by comparing the reconstructed break-in faces with the original break-in samples show that break-in faces are reconstructed with very high quality (purple curve in Figure 1b). Reconstructed break-in faces have better quality compared to reconstructed BRs.

3 Application of *MSK* algorithm to retina point patterns

This study endeavours to tune the *MSK* algorithm on spatial point pattern templates. The attacker has access to the comparison scores issued by the system, and is unable to gain access to the BRs. However, he has access to the break-in set. From the attacker's perspective, a smart method is to try and reconstruct the break-in set templates, first. Then, tune the attack using the knowledge obtained by reconstructing the break-in set. It is expected from the algorithm to successfully reconstruct each break-in template $\mathbf{X}_i,\ 1 \leq i \leq K$ since its corresponding data point $\mathbf{Y}_i,\ 1 \leq i \leq K$ exists in the modelled affine space (having zero distance to itself).

We conducted our experiments over a retinal vascular database, called ESRID (ECG Synchronised Retinal Image Database) [Ha12]. This database is collected by RMIT University and is accessible on request from the authors in [Ha12]. ESRID consists of 414 retinal images of 46 data subjects. Each individual in this dataset has 9 samples of their left eyes. The size of images in ESRID database is 2376×1584 pixels. The retinal point patterns are extracted from spatial graphs that are rescaled and centered on the optic disc. Rescaling sets the fovea on the point $(1,0)$ and the point patterns do not require further registration. The feature points from each image are extracted as real-valued spatial coordinates (x,y), and values can be negative. In experiments that reconstruct the break-in set, templates from the first data subject are considered as references, and the remaining 405 templates constitute the break-in set. In experiments that reconstruct BRs, we performed 46 experiments to reconstruct the BRs that are independent from the break-in set. Each of these 46 experiments reconstructs templates of one data subject using templates from every other data subject as break-in set.

We were interested in investigating the impact of the underlying comparison function on the performance of *MSK* reconstructing point pattern templates. We applied two different point pattern comparison functions: ICP (Iterative Closest Point) and MHD (Modified Hausdorff Distance) [DJ94] in our experiments.

3.1 Adaption of the spatial coordinates format

Here each $\mathbf{X}_i$ is a list of (x,y) coordinates of the points in the break-in sample X_i and would be read as a column vector with x-coordinates followed by y-coordinates. However, the sizes of $\mathbf{X}_i$s vary since different break-in samples have different numbers of points. The attacker needs to modify $\mathbf{X}_1,\ldots,\mathbf{X}_K$ to have the same dimensions by padding each $\mathbf{X}_i$ with enough $(0,0)$ points to increase its length to the maximum length template. After padding, each template will be of length $1,092$. Using *MSK*, any reconstructed templates will have length $1,092$ and a cluster of points close to $(0,0)$ caused by reconstructing the added $(0,0)$ coordinates.

To thoroughly study the impact of this introduced noise, we first focused on break-in samples, which an attacker should be able to reconstruct well (as they can for face images). The idea is that the attacker is well aware of the sizes of the break-in templates. Thus,

when $\mathbf{X}_z$ is generated for a break-in template, the attacker can truncate exactly the auxiliary points it contains. Figure 2 shows the scores obtained by comparing the reconstructed break-in spatial coordinates with the original break-in spatial coordinates both with noise (solid purple curve) and without noise (dashed purple curve) for MHD and ICP.

To estimate the effect of the added noise on reconstruction of BRs we checked the performance of the attack after discarding the points (x,y) from the reconstructed templates, where $-\theta \leq x,y \leq \theta$ for $\theta \in \{0,0.25,0.5,0.75,1\}$. Since the performance of the attack using these values was not markedly different, we only present here the results for $\theta = 0.5$. The blue curves in Figure 2 represent the results of these experiments. Note that the plotting function in Figure 2a over-smooths the curve for reconstructed BRs. The smallest MHD score obtained by comparing the reconstructed BR templates and their corresponding targeted templates equals 2.782. This score is greater than the smallest score for non-mated comparisons which is 2.408.

In Figures 2, the distribution curves for scores obtained by comparing the reconstructed templates with the original templates illustrate that when treating retinal point patterns as spatial coordinates, the reconstructed templates would not be accepted by the system as mated templates. In fact, even discarding all noise introduced to the break-in templates (the dashed purple curves) does not improve the performance of the attack enough to reconstruct templates that could be considered by the comparator as mated templates.

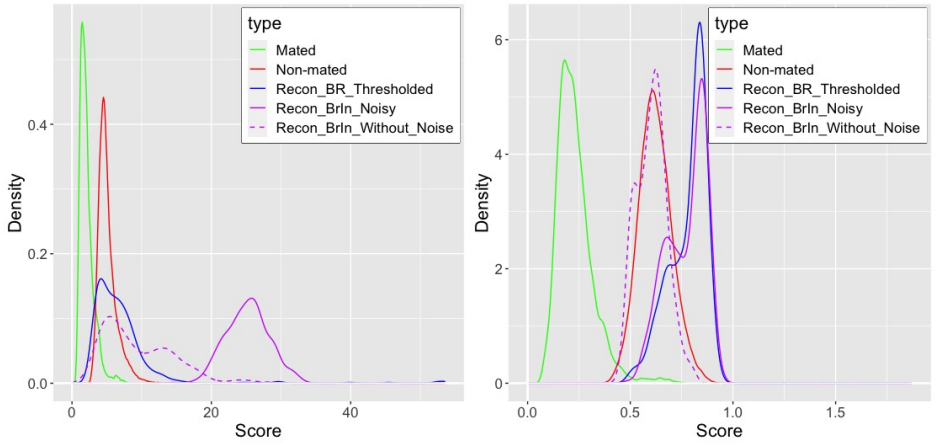

(a) (b)

Fig. 2. Score distributions for mated (green), non-mated (red), and *MSK*-reconstructed templates (blue, with ESRID BRs; solid purple, with ESRID break-in set with noise; dashed purple, with ESRID break-in set noise-reduced) using MHD (a) and ICP (b).

3.2 Adaption of the binary image format

Here, we treat the templates as binary images with discretised intensity values equal to either 255, if a pixel contains a feature point, or 0, otherwise. Feature points in these tem-

plates are real-valued spatial coordinates (x,y). To identify the pixels that contain feature points, the original spatial coordinates were shifted, scaled and rounded to ensure that all the coordinates were positive integers. To minimise rounding error, we applied the inverse transformation to the rounded coordinates and calculated the MHD and ICP distance between the inverted values and the original (x,y) coordinates, with scaling values $1,5,10$ and 15. For accuracy and speed, we selected $s = 10$, resulting in templates of length $1,254,400$.

After performing the attack on a target sample X_t, the returned intensity values in the vector $\mathbf{X}_z$ are greyscale, and must be binarised by thresholding. An attacker can find threshold(s) θ for the reconstructed break-in set then use this information to binarise reconstructed BRs. Following this approach we found that each greyscale break-in reconstruction had its own threshold. If $\theta = 128$, no binarised break-in template is reconstructed. Optimising over the thresholds we found for each break-in reconstruction gave $\theta = 25$ and $\theta = 38$, respectively, for MHD-based and ICP-based systems. However, not all binarised break-in reconstructions with these uniform thresholds will contain feature points. The purple curves in Figure 3 illustrate the performance of the reconstructed break-in set using the optimum thresholds.

Then, we applied those values as thresholds to reconstruct BRs. With $\theta = 25$ (resp. $\theta = 38$) for the MHD-based (resp. ICP-based) system, 32.6% (resp. 28.51%) of the binarised reconstructed BR templates had no feature points in them. We discarded their scores when plotting performance. The blue curves in Figure 3 illustrate the performance of reconstructing BRs using the optimum thresholds.

We see that when treating retinal point patterns as binarised images, the reconstructed BR templates would not be accepted by the system as mated templates. In fact, many of the break-in templates cannot be reconstructed well enough to be considered by the comparator as mated.

4 Conclusion

We applied the *MSK* algorithm to two formats of retina point pattern templates in this paper. The experimental results showed that the performance of this attack on retinal vascular point patterns is not comparable with that on face images. From our point of view, the approximation procedures in calculating $\mathbf{Y}_i$ s and $\mathbf{A}^\dagger$ skew the the values more than a point pattern comparison algorithm can tolerate. Image-based face comparison algorithms can tolerate this amount of deviation since there is a high amount of continuity among their values, whereas for point patterns this is not the case. The performance of this attack on retinal vascular point patterns is so poor that it might not be considered as a threat to privacy and security of the users and their templates. The results using our proposed formats showed that the attack is not even successful in reconstructing break-in templates. For our future work, we would like to check the performance of the *MSK* attack on the point patterns extracted from other vascular biometric characteristics that have reasonable recognition accuracy. We are also interested in exploring the reconstruction of point pat-

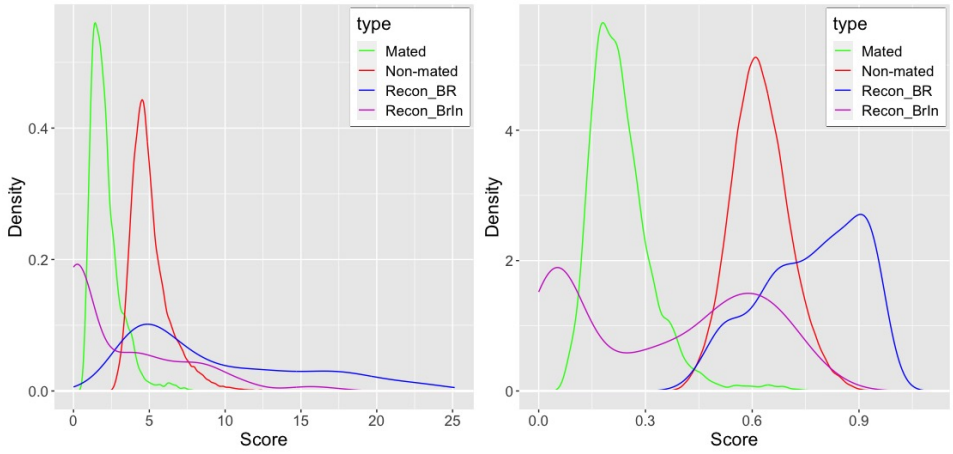

(a) (b)

Fig. 3. Score distributions for mated (green), non-mated (red), and *MSK*-reconstructed templates (blue, with ESRID BRs; purple, with ESRID break-in set) using MHD (a) and ICP (b).

tern and graph-based vascular templates using a hybrid reconstruction method that applies the affine-based reconstruction algorithm followed by a hill-climbing attack.

Acknowledgements The first author was supported by an RMIT University RD Gibson Grant. We thank the reviewers for their comments, the authors in [MSK06] for providing us with their codes for reconstructing faces, and NCI Australia for computational resources. Portions of the research in this paper use the FERET database of facial images collected under the FERET program, sponsored by the DOD Counterdrug Technology Development Program Office.

References

[Ad04] Adler, Andy: Images can be regenerated from quantized biometric match score data. In: Canadian Conference on Electrical and Computer Engineering 2004 (IEEE Cat. No. 04CH37513). volume 1. IEEE, pp. 469–472, 2004.

[Ar16] Arakala, Arathi; Davis, Stephen A; Hao, Hao; Horadam, Kathy J: Value of graph topology in vascular biometrics. IET Biometrics, 6(2):117–125, 2016.

[Be03] Beveridge, Ross; Bolme, David; Teixeira, Marcio; Draper, Bruce: The CSU face identification evaluation system userâs guide: version 5.0. Computer Science Department, Colorado State University, 2(3):1–29, 2003.

[DJ94] Dubuisson, M-P; Jain, Anil K: A modified Hausdorff distance for object matching. In: Proceedings of 12th international conference on pattern recognition. volume 1. IEEE, pp. 566–568, 1994.

[EC16] European Council, General Data Protection: Regulation of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46. Official Journal of the European Union (OJ), 59(1-88):294, 2016.

[Ga13] Galbally, Javier; Ross, Arun; Gomez-Barrero, Marta; Fierrez, Julian; Ortega-Garcia, Javier: Iris image reconstruction from binary templates: An efficient probabilistic approach based on genetic algorithms. Computer Vision and Image Understanding, 117(10):1512–1525, 2013.

[GBG20] Gomez-Barrero, Marta; Galbally, Javier: Reversing the irreversible: A survey on inverse biometrics. Computers & Security, 90:101700, 2020.

[Go14] Gomez-Barrero, Marta; Galbally, Javier; Morales, Aythami; Ferrer, Miguel A; Fierrez, Julian; Ortega-Garcia, Javier: A novel hand reconstruction approach and its application to vulnerability assessment. Information Sciences, 268:103–121, 2014.

[Ha12] Hao, Hao; Sasongko, Muhammad B; Wong, Tien Y; Azemin, Mohd Zulfaezal Che; Aliahmad, Behzad; Hodgson, Lauren; Kawasaki, Ryo; Cheung, Carol Y; Wang, Jie Jin; Kumar, Dinesh K: Does retinal vascular geometry vary with cardiac cycle? Investigative ophthalmology & visual science, 53(9):5799–5805, 2012.

[IS11] ISO/IEC JTC1 SC27 IT Security Techniques: International Standard on Biometric Information Protection. Standard ISO/IEC24745, International Organization for Standardization, 2011.

[MSK06] Mohanty, Pranab; Sarkar, Sudeep; Kasturi, Rangachar: Privacy & security issues related to match scores. In: 2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06). IEEE, pp. 162–162, 2006.

[MSK12] Mohanty, Pranab; Sarkar, Sudeep; Kasturi, Rangachar: , Reconstruction of biometric image templates using match scores, April 24 2012. US Patent 8,165,352.

[Ph98] Phillips, P Jonathon; Wechsler, Harry; Huang, Jeffery; Rauss, Patrick J: The FERET database and evaluation procedure for face-recognition algorithms. Image and vision computing, 16(5):295–306, 1998.

[Ph00] Phillips, P Jonathon; Moon, Hyeonjoon; Rizvi, Syed A; Rauss, Patrick J: The FERET evaluation methodology for face-recognition algorithms. IEEE Transactions on pattern analysis and machine intelligence, 22(10):1090–1104, 2000.

[TP91] Turk, Matthew; Pentland, Alex: Face recognition using eigenfaces. In: Proceedings. 1991 IEEE computer society conference on computer vision and pattern recognition. pp. 586–587, 1991.