Deep Multi-Finger Fuzzy Commitment

Hans Geißner, Valentina Fohr and Christian Rathgeb
da/sec – Biometrics and Security Research Group
Hochschule Darmstadt, Germany
hans.geissner@h-da.de

***Abstract*—Biometric cryptosystems like the fuzzy commitment scheme are designed to protect sensitive biometric data through error tolerant retrieval of keys. By binding biometric data with cryptographic keys, the fuzzy commitment scheme facilitates privacy-preserving biometric verification, ensuring that biometric data remains protected even in the event of a system breach. This paper investigates the application of the fuzzy commitment scheme to a multi-instance biometric system using deep learning-based feature extractors for fingerprints. By leveraging multiple fingerprint instances, the scheme enhances both security in terms of in False Accept Security (FAS) and recognition performance in terms of Genuine Match Rate (GMR). Notably, a four-finger configuration achieves a near-perfect GMR of 99.35% and a FAS of 17.7 bits, surpassing single- and two-finger configurations. These findings highlight the potential of multi-instance fusion for achieving high levels of security as well as usability in biometric cryptosystems while addressing challenges in error correction.**

***Index Terms*—Biometric cryptosystems, fingerprint recognition, fuzzy commitment scheme, deep neural networks**

I. Introduction

STORING biometric reference data, i.e. templates, for authentication purposes poses a unique challenge: biometric data is sensitive and immutable, raising significant privacy concerns if compromised, while traditional cryptographic methods like hashing cannot be directly applied due to the inherent variability of biometric traits. *Biometric Template Protections (BTPs)* [1, 2] addresses these challenges by adhering to the requirements of biometric information protection, facilitating privacy-preserving storage and accurate comparison of biometric data. These requirements, as defined in ISO/IEC IS 24745 [3], include:

Irreversibility: It should be infeasible to reconstruct biometric data from the protected template.

Unlinkability: It should be infeasible to determine if two protected template correspond to the same individual.

Revocability and Renewability: It should be possible to issue a new protected template without revealing additional information.

Performance Preservation: The recognition accuracy of the protected system should be comparable to that of unprotected systems.

This work was supported in part by the German Federal Ministry of Education and Research; in part by the Hessian Ministry of Higher Education, Research, Science and the Arts within their joint support of the National Research Center for Applied Cybersecurity (ATHENE).

The authors are with the da/sec—Biometrics and Security Research Group, Hochschule Darmstadt, 64295 Darmstadt, Germany (e-mail: hans.geissner@h-da.de; valentina.fohr@stud.h-da.de; christian.rathgeb@h-da.de)

Biometric Cryptosystems (BCSs), a category of BTP schemes, protect biometric templates by combining biometric data with a cryptographic key. Instead of directly comparing biometric samples, these systems verify the validity of an associated key, enabling indirect biometric authentication. The *fuzzy commitment* scheme, is one of the most widely used BCSs, alongside the fuzzy vault scheme. The scheme protects a biometric template by using it as noise to conceal a cryptographic key, which is encoded by an Error Correction Code (ECC). The key can potentially be recovered by removing the noise using a biometric probe that is sufficiently similar (i.e., with only a correctable amount of errors) to the original template. Since the binding between the biometric data and the key can be reversed if the key is exposed, revealing the original biometric data, the security of the system depends not only on the robustness of the scheme but also on the security of the key itself [4]. Therefore, the scheme should not only have an adequate security level but also a sufficiently long key.

Scientific research has shown that the security levels of biometric cryptosystems based on a single source of biometric data are significantly lower than the security standards deemed secure in traditional cryptography. As a result, several studies have proposed multi-biometric template protection systems [5] which fuse biometric data from multiple sources, e.g., multiple characteristics or multiple instances of a single characteristic, before applying template protection. This approach increases the usable information within the system, thereby improving both the accuracy and security of the BCS. However, security is only enhanced when biometric fusion takes place at the feature level [6], similar to using a single long password instead of several short ones. In the case where multiple characteristics are fused, additional challenges related to imbalances in weighting due to varying feature lengths and biometric variance between different characteristics arise [7]. In such systems, attackers may exploit the characteristic with the weakest performance or the largest weight to launch false accept attacks. By focusing on the fusion of multiple instances of the same biometric characteristic, such as combining several fingerprints from the same individual, these vulnerabilities can be mitigated as the same feature extractors can be used for each instance. With recent advancements in deep learning, robust methods for extracting fixed-length real-valued feature vectors from biometric characteristics, including fingerprints, have been introduced. As the fuzzy commitment scheme requires biometric templates to be binary vectors, these deep learning-based feature extractors are well-suited, since their

feature vectors can be easily binarized. This study aims to leverage the effectiveness of the fuzzy commitment scheme in protecting deep learning-based biometric templates and demonstrate how multi-instance fusion can be used to increase the performance and security of the fuzzy commitment scheme. For this purpose, a case study using multiple instances of fingerprints is conducted using a state-of-the-art deep neural network for feature extraction.

The remainder of this work is organized as follows: In section II, the background and related work are presented. In section III, the methodology is described. In section IV, the experimental results are presented. Finally, in section V, future work is discussed, and the conclusion is provided.

II. Related Work

A. The Fuzzy Commitment Scheme

The fuzzy commitment scheme, introduced by Juels and Wattenberg [8], protects biometric templates while allowing error-tolerant verification. It binds a fixed-length binary vector to a cryptographic key, that is encoded using an ECC that detects and corrects errors within a binary vector. To validate the correctness of a retrieved key, a hash of the key is additionally stored. The binding is achieved by applying a bitwise Exclusive-OR (XOR) operation between the encoded key and the biometric template where both are required to be of same length. Since the XOR operation is self-inverse, applying another bitwise XOR with a biometric probe, exposes an erroneous version of the encoded key, where the differences between the original template and the probe correlate with the errors. The original key can be recovered depending on the number of errors, their distribution, and the chosen error correction code. The correctness of a retrieved key is verified by computing its hash and comparing it to the stored hash of the original key. Notably, once the correct key is retrieved, the original template is also revealed.

Early works have applied the scheme to binary iris codes [9]. To apply the scheme to characteristics, for which the features are not commonly extracted in the form of binary vectors, feature type transformations have been used to make them compatible with the scheme. For instance, this concept was applied to fingerprints in [10] and voice in [11]. More recently, deep learning has been used to extract feature vectors from facial data, which are then binarised and protected using the fuzzy commitment scheme [12].

It has been demonstrated that a statistical attack can compromise the fuzzy commitment scheme in offline attack scenarios, especially when the key is distributed across multiple blocks, exposing the key as well as the original template [13, 14]. This attack operates by performing multiple non-mated comparisons, constructing a histogram of the decoded codewords for each block, and identifying the most frequently occurring codewords to reconstruct a key. If the average Hamming distance between two uncorrelated (i.e., non-mated) template is significantly below 0.5, the correct codeword for each block can be identified with a relatively small number of non-mated comparisons, ultimately exposing the correct key. Moreover, a linkage attack has been demonstrated that is able to determine if two protected templates correspond to the same individual, due to the linear property of the used error correcting codes [15]. It was shown that this attack can be mitigated by performing a public pseudo-random permutation that shuffles the individual bits before applying the fuzzy commitment scheme [10, 16].

Recent studies using deep learning-based feature extraction have shown that performance levels sufficient for traditional biometric systems may be inadequate for fuzzy commitment scheme, as a single false accept can enable the recovery of the committed binary template. For face recognition, such attacks enable approximate reconstruction of the facial image [17]. For gait recognition, the effective security against false accepts was estimated to be comparable to a 4-digit PIN [18].

TABLE I
Overview of most relevant works on multi-biometric fuzzy commitment schemes.

Ref.	Year	Characteristic(s)	Dataset(s)	FNMR (in %)	FMR (in %)	Key size (in bits)
[19]	2009	3D Face	FRGC	$\sim 22\%$	0.25%	155 bits
[20]	2011	Two irises	CASIA-v3	5.56%	0.01%	128 bits
[21]	2012	Fingerprint Iris, Face	FVC02 DB2, CASIA-v1, XM2VTS	$\sim 1\%$	$\sim 0.0\%$	n.a.
[22]	2013	Face, Iris	NIST-ICE, FRGC	0.89%	0.0%	217 bits

B. Multi-Biometric Fuzzy Commitment Schemes

The above table provides an overview of relevant works on multi-biometric fuzzy commitment schemes, presenting their False Non-Match Rate (FNMR), False Match Rate (FMR), and key sizes. These schemes typically focus on multi-algorithmic, multi-instance, and multi-modal fusion approaches. Kelkboom et al. [19] and Rathgeb et al. [20] explore multi-algorithm and multi-instance fusions at the feature level, where features from different sources are combined and filtered for reliability. Nagar et al. [21] use feature type transformations for fingerprint, face, and iris embeddings, reporting high performance at a security level of 53 bits. Kanade et al. [22] combines multi-instance fusion with multi-modal fusion (face and iris), using a weighted error correction for the imbalance of feature sizes.

While these methods improve error distribution and recognition performance, the key sizes reported are not accurate measures of security, as false matches can reveal keys. Additionally, works reporting an FMR of 0% likely suffer from limited non-mated comparisons. Furthermore, these approaches predate the use of deep learning-based feature extractors in biometric systems.

III. Methodology

A. Feature Extraction and Binarisation

This work explores the viability of multi-instance fusion for securing multiple fingerprint templates using the fuzzy commitment scheme. Deep learning-based feature extractors

are leveraged to generate fingerprint templates suitable for protection. An overview of the system is shown in Figure 1.

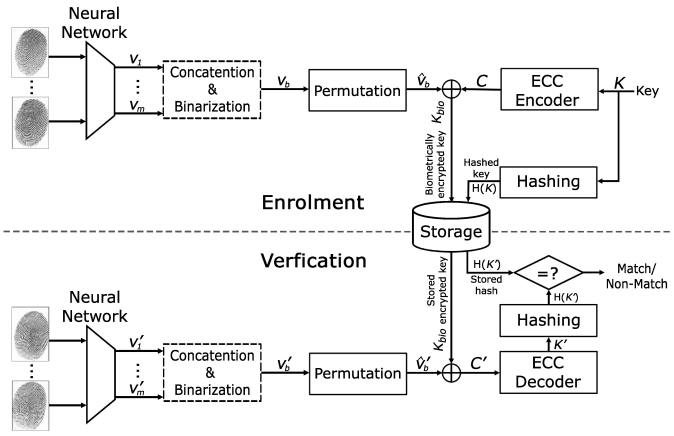

Fig. 1. Overview of the fuzzy commitment scheme for multi-finger biometrics. The diagram illustrates both enrollment and retrieval processes, including feature extraction from multiple fingerprints, concatenation, binarization, and error-correcting code (ECC) encoding and decoding.

In the multi-instance scenario, a single deep learning-based feature extractor can be applied to all fingerprint instances, eliminating the need for separate extractors for each template. Deep learning-based feature extractors produce fixed-length real-valued vectors, typically of size $n = 2^x$ (e.g., 256 or 512), due to their training with differentiable loss functions such as the Euclidean distance and hardware optimization. For a multi-instance system, m feature vectors are generated and concatenated into a single multi-biometric feature vector $v = v_1||\dots||v_m = (x_1, \dots, x_{n\cdot m})$, of size $n \cdot m$. To protect these feature vectors with the fuzzy commitment scheme, the real-valued vectors are binarised into binary strings. Since these vectors typically have distributions centred around zero, their features are approximately evenly split between positive and negative values. This allows binarisation of features according to the sign $q(x) = \begin{cases} 1 \text{ if } x \geq 0 \\ 0 \text{ else} \end{cases}$, mapping positive values to 1 and negative values to 0, ensuring an equal probability of ones and zeros. Using this approach, the vector is binarised by binarising each individual feature, $v_b = (q(x_1), \dots, q(x_{n\cdot m}))$. While binarisation methods that map features to more than one bit can lead to better performance, as shown in Drozdowski et al. [23], these methods need to be trained on the respective feature extraction methods. Another advantage of applying this simple binarisation method to deep learning-based feature vectors is that, since each binarised feature is expected to have an equal probability of being zero or one, two uncorrelated vectors v and v' (i.e., non-mated feature vectors) will, on average, have a Hamming distance of 0.5. This is expected to mitigate the statistical attack proposed in [14].

B. Enrolment and Verfication

A subject is enrolled within the scheme by collecting a biometric reference, i.e. fingerprint(s), from the subject and processing it as described in subsection III-A. A secret key κ is chosen and encoded using an ECC, which results in $c = enc(\kappa)$. The key-size k depends on the choice of the ECC, as the size of the encoded key needs to be equal to the template size $n \cdot m$. Moreover, the hash $H(\kappa)$ is calculated, stored, and used afterwards to verify the correctness of a retrieved key. Additionally, the hash is used to seed a public bijection σ that is used to shuffle the features in the feature vector $\hat{v}_b = \sigma(v_b)$. This step serves a role similar to that of cryptographic salting, introducing pseudo-randomness into the stored record. Specifically, it prevents different templates of the same subject from being directly correlated with each other [16]. The main operation of the fuzzy commitment scheme is applying a bitwise XOR between the biometric template and the encoded key, creating the binding $\delta = c \oplus \hat{v}_b$. The pair $(\delta, H(\kappa))$, called the commitment, is stored.

During verification, a biometric probe is collected from the subject, and the corresponding feature vector v' is used to obtain v'_b using the same binarisation process as in enrolment. Using $H(\kappa)$, the public bijection σ is reconstructed, and the biometric probe is permuted to obtain $\hat{v}'_b = \sigma(v'_b)$. Next, the bitwise XOR operation is performed between the permuted biometric probe and the binding, i.e., $\delta \oplus \hat{v}_b$. The resulting vector is then fed into the decoder function of the ECC to compute a candidate key $\kappa' = dec(\delta \oplus \hat{v}_b)$. Finally, the correctness of the candidate key is verified by computing its hash and comparing it to the stored hash, i.e., $H(\kappa') \stackrel{?}{=} H(\kappa)$. The key can be correctly retrieved if the distance vector $\hat{v}_b \oplus \hat{v}'_b$ between the biometric reference and probe satisfies the error-correction capability of the code, meaning $|\hat{v}_b \oplus \hat{v}'_b|$ is within the allowed threshold, where $|\cdot|$ denoted the Hamming weight.

C. Error Correction Codes

The fuzzy commitment scheme uses ECCs to handle intra-class variance in biometric data. Small differences between the biometric template used during key binding and retrieval are treated as errors that the ECC corrects. Choosing the right configuration ensures successful retrieval while sustaining a low probability of false matches. The number of correctable errors cannot be freely chosen but is implicitly determined by the used ECC. Additionally, in some cases (e.g., multi-level codes), the number of correctable errors may not be fixed and depends on the distribution of errors. Bit-level error correction is ideal for biometric data in binary form when single-bit errors are expected. Hadamard codes, commonly used for small bit-blocks, encode bit strings into larger code words with minimum pairwise distances, allowing errors to be corrected. However, if errors exceed this threshold, the correction fails, i.e. the string may be mapped to an incorrect code word. In the proposed system, the key κ consisting of k bits is encoded to a string of $n \cdot m$ bits, matching size of the binary feature vector v. Hadamard codes encode k bits into 2^{k-1} bits and can correct up to $2^{k-2} - 1$ errors. To retrieve longer keys, the key is divided into b blocks, each encoded separately, resulting in $k = b\left(\log_2\left(\frac{n}{b}\right) + 1\right)$. This configuration fails if any block exceeds its error correction limit, $\frac{2^{n\cdot m}}{4\cdot b} - 1$. The deployed random permutation can mitigate this issue to some extent, by redistributing errors across the feature vector

[20]. While this reduces the risk of burst errors, block-level errors can still occur. To address this, Reed-Solomon codes are additionally applied. These codes treat blocks as finite-field elements and correct errors by adding redundancy. Keys are segmented, encoded with Reed-Solomon codes, and further encoded with Hadamard codes. During retrieval, minor bit flips are corrected in a first step, followed by burst error correction in the second step. This multi-level ECCs enables iterative decoding to enhance correction capabilities across blocks [24].

D. Experimental Setup

Experiments are conducted using the MCYT-330 fingerprint database [25], which contains samples from each finger of 330 subjects. In this study, each hand of a subject is considered as a separate subject, using all fingers except the thumb. This approach results in 660 subjects, each comprising four fingers. For feature extraction, an re-implementation of the DeepPrint method is employed, generating fixed-length real-valued feature vectors of size 512 [26]. Experiments are performed by combining features from one, two, and four fingers in the fuzzy commitment scheme, varying the setups of the ECC. Each experiment includes 3,960 mated comparisons and 2,609,640 non-mated comparisons. Biometric performance is evaluated using Genuine Match Rate (GMR) and FMR. The GMR represents the rate at which mated comparisons are correctly classified as matches, while the FMR represents the rate at which non-mated comparisons are incorrectly classified as matches. The False Accept Security (FAS) is used to assess system security because false accept attacks tend to be the most effective once statistical and linkage attacks have been ruled out, and they are generally more potent than brute-force attacks. In such attacks, an adversary simulates non-mated authentication attempts to cause a false match, which ultimately reveals the key and the protected biometric data. This assumes the attacker has full knowledge of the algorithm, access to protected templates, and statistical data about the biometric feature vectors. The FAS represents the expected number of non-mated comparisons needed to produce a false match and is expressed in bits as $\log_2(\text{FMR}^{-1})$ [27].

IV. Results and Discussion

A. Results

The experimental results are summarized in Table II for one, two, and four fingers. Since the deployed permutation can affect the outcome of a comparison trial, the experiments are repeated ten times and the results are averaged. As the experimental results were consistent with only negligible variance, these variations are not further discussed. Notably, for the one-finger system, the average Hamming distance for non-mated comparisons was found to be slightly lower than the expected 0.5, averaging 0.4928. Although this deviation might theoretically make statistical attacks more promising than brute-force attacks, we still expect the false accept attack to be more effective.

Using one finger, the scheme reaches a FAS of 12.06 bits at a GMR of 95.91%, with a key-size of 72 bits. In a naive configuration where there is only a single block, the system achieves a GMR of 99.39 and a FAS of 6.78 bits. The system using two fingers reaches a FAS of 16.38 and a GMR of 97.45% at a key-size of 112. Notably, the naive configuration improves the FAS to 9.06 bits while also improving the GMR slightly to 99.87%. The system based on four fingers reaches a FAS of 17.7 bits with a near-perfect GMR of 99.35% at a key-size of 112 bits. While this marks only a slight improvement of approximately 1 bit in terms of security, the improved GMR suggests that false non matches are four times less likely, meaning that while the four finger system is only slightly more secure than the two-finger system, the convenience is increased, since capturing four fingers or two fingers makes practically no difference in terms of convenience, but the frequency of false no matches and therefore of required recaptures is significantly reduced. For this setup, the naive configuration reaches a FAS of 12.37 bits at a perfect GMR of 100%.

Notably, there is a strong correlation between the key-size and the FAS. This can be attributed to the fact that in an ECC, a larger message length (key-size) is generally associated with a lower codeword distance and therefore with a lower number of correctable errors. This means that for larger key-sizes, the similarity of a comparison needs to be higher to produce a match. However, while the same principle is true for the GMR, meaning that the GMR should be lowered for higher k, there are more deviations. For example, for one finger with a key size of 56 bits, the GMR is only 94.77%, whereas at a key size of 60 bits, the GMR is 97.30%. While smaller key sizes generally allow for a higher maximum number of correctable errors, the use of block-level error correction increases the minimum number of correctable errors. This demonstrates the effectiveness of two-step ECC in improving performance but also highlights how variability in error correction capability can negatively impact system performance.

B. Discussion

The findings confirm that the multi-instance fusion effectively increases both system performance and security. Notably, the four-finger setting demonstrates improvements in both GMR and FAS compared to single- and two-finger configurations. Compared to the works presented in Table I, the four-finger setup provides similar performance levels, while also offering a non-zero value for FMR, which surpasses the performance of other works reporting non-zero FMRs. Regarding key-size, our work falls at the lower end, with a maximum key size of 128 bits. However, since fuzzy commitment schemes are likely to yield a lower FAS, this becomes less significant.

In contrast to more complex approaches such as homomorphic encryption, the fuzzy commitment scheme offers practical advantages in terms of runtime. The underlying error-correcting codes are computationally lightweight, relying on simple algebraic operations, which allows efficient implementation even in constrained environments.

TABLE II
RESULTS FOR 1-, 2-, AND 4-FINGER SETUPS. THE 4-FINGER CONFIGURATION ACHIEVES THE BEST TRADEOFF BETWEEN GMR AND FAS.

One Finger					Two Fingers					Four Fingers				
k	b	GMR%	FMR%	FAS	k	b	GMR%	FMR%	FAS	k	b	GMR%	FMR%	FAS
10	1	99.39	0.9124	6.78	11	1	99.87	0.1877	9.06	12	1	100.00	0.0189	12.37
18	2	98.61	0.4031	7.95	20	2	99.83	0.1162	9.75	22	2	99.98	0.0107	13.19
32	4	97.85	0.1675	9.22	36	4	99.69	0.0506	10.95	40	4	99.90	0.0044	14.49
16	4	98.89	0.5449	7.52	18	4	99.85	0.1307	9.58	20	4	99.99	0.0131	12.90
56	8	94.77	0.0349	11.48	64	8	99.30	0.0116	13.08	72	8	99.76	0.0012	16.33
42	8	97.90	0.1261	9.63	48	8	99.66	0.0354	11.47	54	8	99.84	0.0031	14.98
28	8	98.61	0.3016	8.37	32	8	99.79	0.0744	10.39	36	8	99.94	0.0067	13.86
96	16	83.08	0.0021	15.51	112	16	97.45	0.0012	16.38	128	16	99.35	0.0005	17.70
84	16	92.95	0.0093	13.40	98	16	98.92	0.0034	14.84	112	16	99.67	0.0006	17.40
72	16	95.91	0.0235	12.06	84	16	99.27	0.0081	13.60	96	16	99.73	0.0009	16.75
60	16	97.30	0.0564	10.79	70	16	99.49	0.0167	12.55	80	16	99.80	0.0016	15.91
48	16	98.13	0.1119	9.80	56	16	99.65	0.0311	11.65	64	16	99.86	0.0029	15.09

V. CONCLUSION

This study explores the application of the fuzzy commitment scheme in a multi-instance biometric system, utilizing deep learning-based feature extractors for fingerprint templates. The work presents a system designed to meet the requirements of Biometric Template Protection. Irreversibility is ensured by the binding process of the fuzzy commitment scheme and statistical attacks are mitigated through the use of deep learning-based feature extractors. Unlinkability is achieved by applying (pseudo)random permutations, and renewability is implicitly fulfilled by satisfying both irreversibility and unlinkability. While an in-depth evaluation of performance preservation was not conducted, our results indicate that the system maintains competitive recognition accuracy. Notably, the four-finger configuration achieved a near-perfect GMR of 99.35% and a FAS of 17.7 bits with 128-bit key bound. However, as discussed, the effectiveness of multilevel ECC depends on the error distribution, meaning the number of correctable errors is not constant. This variability can impact both performance and security. Therefore, future work should explore the integration of additional fusion techniques, such as interleaving, alongside the deployed permutation, to optimize error distributions and stabilize the error-correcting capabilities of the code. Another direction for future work could be to systematically evaluate whether deep learning-based feature extractors can effectively mitigate statistical attacks.

REFERENCES

[1] M. Barni, G. Droandi, and R. Lazzeretti, "Privacy protection in biometric-based recognition systems: A marriage between cryptography and signal processing," *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 66–76, 2015.

[2] P. Melzi, C. Rathgeb, R. Tolosana, R. Vera, and C. Busch, "An overview of privacy-enhancing technologies in biometric recognition," *ACM Computing Surveys (CSUR)*, Mai 2024.

[3] ISO/IEC JTC1 SC27 Security Techniques, *ISO/IEC 24745:2022. Information Technology - Security Techniques - Biometric Information Protection*, International Organization for Standardization, 2022.

[4] T. Ignatenko and F. M. J. Willems, "Information leakage in fuzzy commitment schemes," *IEEE Transactions on Information Forensics and Security*, vol. 5, no. 2, pp. 337–348, 2010.

[5] C. Rathgeb and C. Busch, "Multi-biometric template protection: Issues and challenges," in *New Trends and Developments in Biometrics*, J. Yang and S. J. Xie, Eds. Rijeka: IntechOpen, 2012, ch. 8. [Online]. Available: https://doi.org/10.5772/52152

[6] J. Merkle, T. Kevenaar, and U. Korte, "Multi-modal and multi-instance fusion for biometric cryptosystems," in *Intl. Conf. of Biometrics Special Interest Group (BIOSIG)*, 2012, pp. 1–6.

[7] C. Rathgeb, B. Tams, J. Merkle, V. Nesterowicz, U. Korte, and M. Neu, "Multi-biometric fuzzy vault based on face and fingerprints," in *Proc. Intl. Joint Conf. on Biometrics (IJCB)*, 2023.

[8] A. Juels and M. Wattenberg, "A fuzzy commitment scheme," in *6th ACM Conf. on Computer and Communications Security*, 1999, pp. 28–36.

[9] F. Hao, R. Anderson, and J. Daugman, "Combining crypto with biometrics effectively," *IEEE Computer*, vol. 55, pp. 1081–1088, 2006.

[10] E. J. C. Kelkboom, J. Breebaart, T. A. M. Kevenaar, I. Buhan, and R. N. J. Veldhuis, "Preventing the decodability attack based cross-matching in a fuzzy commitment scheme," *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 1, pp. 107–121, 2011.

[11] S. Billeb, C. Rathgeb, H. Reininger, K. Kasper, and C. Busch, "Biometric template protection for speaker recognition based on universal background models," *IET Biometrics*, vol. 4, no. 2, pp. 116–126, 2015.

[12] B. P. Gilkalaye, A. Rattani, and R. Derakhshani, "Euclidean-distance based fuzzy commitment scheme for biometric template security," in *2019 7th International Workshop on Biometrics and Forensics (IWBF)*, 2019, pp. 1–6.

[13] A. Stoianov, T. Kevenaar, and M. van der Veen, "Security issues of biometric encryption," in *2009 IEEE Toronto International Conference Science and Technology for Humanity (TIC-STH)*, 2009, pp. 34–39.

[14] C. Rathgeb and A. Uhl, "Statistical attack against fuzzy commitment scheme," *IET Biometrics*, vol. 1, no. 2, pp. 94–104, 2012.

[15] K. Simoens, P. Tuyls, and B. Preneel, "Privacy weaknesses in biometric sketches," in *2009 30th IEEE Symposium on Security and Privacy (S&P)*, 2009, pp. 188–203.

[16] B. Tams, "Decodability attack against the fuzzy commitment scheme with public feature transforms," 2014. [Online]. Available: https://arxiv.org/abs/1406.1154

[17] D. Keller, M. Osadchy, and O. Dunkelman, "Fuzzy commitments offer insufficient protection to biometric templates produced by deep learning," *CoRR*, vol. abs/2012.13293, 2020.

[18] T. V. hamme, E. Rúa, D. Preuveneers, and W. Joosen, "On the security of biometrics and fuzzy commitment cryptosystems: A study on gait authentication," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 5211–5224, 2021.

[19] E. Kelkboom, X. Zhou, J. Breebart, R. Veldhuis, and C. Busch, "Multi-algorithm fusion with template protection," in *Proc. 3rd Intl. Conf. on Biometrics: Theory, Applications and Systems (BTAS)*. IEEE Computer Society, September 2009, pp. 1–8.

[20] C. Rathgeb, A. Uhl, and P. Wild, "Reliability-balanced feature level fusion for fuzzy commitment scheme," in *2011 International Joint Conference on Biometrics (IJCB)*, 2011, pp. 1–7.

[21] A. Nagar, K. Nandakumar, and A. Jain, "Multibiometric cryptosystems based on feature-level fusion," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 1, pp. 255–268, 2012.

[22] S. Kanade, D. Petrovska-Delacrétaz, and B. Dorizzi, *Obtaining Cryptographic Keys Using Multi-biometrics*. London: Springer London, 2013, pp. 123–148.

[23] P. Drozdowski, F. Struck, C. Rathgeb, and C. Busch, "Benchmarking binarisation schemes for deep face templates," in *Intl. Conf. on Image Processing (ICIP)*. IEEE, October 2018, pp. 191–195.

[24] J. Bringer, H. Chabanne, G. Cohen, B. Kindarji, and G. Zémor, "Theoretical and practical boundaries of binary secure sketches," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 4, pp. 673–683, 2008.

[25] J. Ortega-Garcia, J. Fierrez-Aguilar, D. Simon, J. Gonzalez, M. Faundez-Zanuy *et al.*, "MCYT baseline corpus: a bimodal biometric database," *Vision, Image and Signal Processing, IEEE Proc.*, vol. 150, no. 6, pp. 395–401, December 2003.

[26] T. Rohwedder, D. Osorio-Roig, C. Rathgeb, and C. Busch, "Benchmarking fixed-length fingerprint representations across different embedding sizes and sensor types," in *Proc. Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*. Gesellschaft für Informatik e.V., September 2023.

[27] C. Rathgeb and A. Uhl, "A survey on biometric cryptosystems and cancelable biometrics," *EURASIP Journal on Information Security*, vol. 3, March 2011.