Significant Feature Based Representation for Template Protection

Deen Dayal Mohan , Nishant Sankaran, Sergey Tulyakov,
Srirangaraj Setlur, Venu Govindaraju,
Department of Computer Science and Engineering, University at Buffalo, Buffalo, New York, USA

Abstract

The security of biometric templates is of paramount importance. Leakage of biometric information may result in loss of private data and can lead to the compromise of the biometric system. Yet, the security of templates is often overlooked in favour of performance. In this paper, we present a plug-and-play framework for creating secure face templates with negligible degradation in the performance of the system. We propose a significant bit based representation which guarantees security in addition to other biometric aspects such as cancelability and reproducibility. In addition to being scalable, the proposed method does not make unrealistic assumptions regarding the pose or illumination of the face images. We provide experimental results on two unconstrained datasets - IJB-A and IJB-C.

1. Introduction

Biometrics based authentication systems have been overwhelmingly successful in providing security to a variety of applications. Biometric authentication is more convenient and secure than other authentication tokens that have to be memorized, such as passwords. Even though adoption of biometric authentication systems is increasing rapidly, a key advantage of information or password-based systems is the ease of cancelability associated with the authentication tokens. In the scenario where the stored token (password) is compromised, it can be changed and a new password can be registered. Since biometric data is inherently non-cancelable there is a need to create cancelable representations for biometric data. It is also important for any biometric system to create a representation that would prevent any leak in user information if the system is compromised. This is extremely important in preserving the privacy associated with user data. Considering face recognition systems in particular, [11] discussed how model inversion attacks could potentially expose the weakness of any facial biometric system and lead to leakage of sensitive user information. Analysis of biometric systems also indicates that, regardless of the biometric modality under consideration, the chances of two independent recordings of data (sensor recordings) from the same data source matching exactly is infinitesimally small. In light of these challenges and the concerns presented by the aforementioned work and other similar approaches, there is a clear need for creating more robust and secure biometric systems.

Prior research [18] [32] [25] has focused on addressing some of these challenges. Approaches such as [26] have created a robust representation by allocating maximum entropy code words for each user with no pre-defined correlation with the original biometric modality (the users face). This property makes attacks on the template very difficult, leaving brute force attacks in the code domain and complex dictionary attacks in the input domain as being the only feasible options. Approaches based on a fuzzy commitment scheme [22] have used error-correcting codes to provide exact matching for face templates. Most of the research efforts to address these challenges in facial biometric systems have limitations while adapting to real-world applications. These limitations can be attributed mainly to three factors. 1) The images used for developing and testing these methods are constrained facial images, restricting their adaptability to a large number of real-world applications. Face images in real-world applications often have large variations in pose or illumination and may have partial occlusions. 2) Scalability of such methods is also a major concern as most of the existing frameworks are limited in the number of distinct identities or subjects that they can handle reliably. The number of identities in the datasets used to test these approaches are far fewer than what is typically seen in an average real-world system. 3) Majority of the existing research also assumes that the target identities are known *a priori* when the security framework is built. This assumption typically does not hold in real world systems where new identities are enrolled frequently.

In the current work, we address all the three aforementioned issues by developing a method which creates a robust significant feature based template that reduces information leak and enables exact matching of facial templates. The plug-and-play framework that is developed can be an add-on to any existing face recognition system. We demon-

strate the capability of the framework for handling unconstrained facial images by reporting performance on the IJB-A dataset and demonstrate the scalability of the framework by reporting performance on the IJB-C dataset.

2. Related Work

There is a large volume of literature in the area of biometric security and privacy and we will primarily focus on methods which are related to face template security. Methods that use fuzzy commitment or fuzzy vault schemes [2], [33] have been explored in the past. Fuzzy vault schemes suffered from the drawback that majority of the data was stored in the open and required substantial storage capabilities. Multiple works [7] [24] [30] have also explored the use of user-specific inputs such as a password along with the biometric data. [31] [19] proposed creating user-specific random projections which when applied on biometric data, obfuscate the information. [25] used local facial region based hashing to improve the security of facial biometrics. [26] improved this method by creating a mapping of user-specific biometric data to maximum entropy binary bits thereby separating biometric information from the information being stored. The binary code assigned to each user is hashed using the cryptographic hash function(SHA-512). Thus, the transformed face template is the cryptographic hash of the binary code assigned to the user. [15] created a more robust mapping using a deeper convolution net architecture. [8] introduced deep learning based quantization hashing to improve privacy and prevent information leakage. Chee et al [6] proposed a modification to the winner take all (WTA) hashing method to provide stronger security against ARM attacks (Attack via Record Multiplicity). Methods for transforming features into a new domain so that it is non-invertible, have also been explored. [28] provided three non-invertible transforms, namely cartesian, polar and functional, for generating cancelable face and fingerprint templates. They achieve high template security but the face recognition performance is low. Expanding on these approaches, researchers have also focused on combining the biometric cryptosystems with unidirectional transformation functions. [10] proposed a similar approach to face template security. [29] explored the use of multimodal biometric fusion using deep networks and increasing security using error correcting codes. [18] showed the value of using binary features in increasing security of biometric data.

Face recognition has been a well explored area of computer vision and numerous methods have been proposed in the previous decades [1][4][20][23]. Face recognition systems aim to extract features that help distinguish face images of different people. Recent face recognition research has mainly focused on three approaches. First, creating deep neural networks trained on multiple large datasets such as MS-Celeb-1M [12], UMDFaces[3], VGGFace2[5] which has shown that training on multiple large-scale datasets improves the performance of the system [27][21]. Secondly, applying modern and deeper convolution architectures [13][14] to improve the performance of face recognition. Finally, designing better loss functions to optimize in order to create more discriminative features[9][35]. We present a method to ensure the security of the biometric templates generated by such state-of-the-art face recognizers while minimizing its impact on face recognition performance.

3. Methodology

Having presented the motivation for constructing a more robust and secure representation of biometric templates, we explain our method for creating a scalable, template security framework.

3.1. Feature Representation with Maximum Entropy Bits

Consider the problem of creating a robust representation for facial biometric data. The face recognizer learns a function which maps $F : I \mapsto R^d$ that maximizes the discriminative ability of $f_i \in R^d$ where f_i refers to the feature embedding and d represents the feature dimension. Given two feature embeddings f_i and f_j, the similarity between them is measured using cosine similarity. To enhance separability, during training an L2 normalization layer is applied, similar to [27], which ensures that the magnitude of the feature embeddings f_i are ignored and only the orientations are considered for maximizing the separability of classes. The softmax loss acting on the L2 normalized features, encourages the feature embeddings belonging to a specific identity to be in close proximity to one another on the surface of the d dimensional hypersphere. Better separation is achieved by allowing only the direction of the feature vectors to influence the loss, thereby prompting the recognizer to focus equally on both hard samples (with small feature vector magnitudes) and easy samples (with large feature vector magnitudes). In order to further enhance separability of feature embedding, techniques such as hard negative mining are used along with loss functions such as triplet loss. The application of triplet loss doesn't restrict the feature embedding to be on a hypersphere, rather it spreads it across a d dimensional manifold.

In traditional biometric systems, feature embeddings are generated for a set of face images of users at the time of enrollment. These embeddings (gallery features) are saved in a conventional database system. Storing gallery feature embeddings or templates in the above mentioned manner may result in leakage of sensitive user information if the security of the database is compromised. This is owing to the fact that feature embeddings have a high degree of association

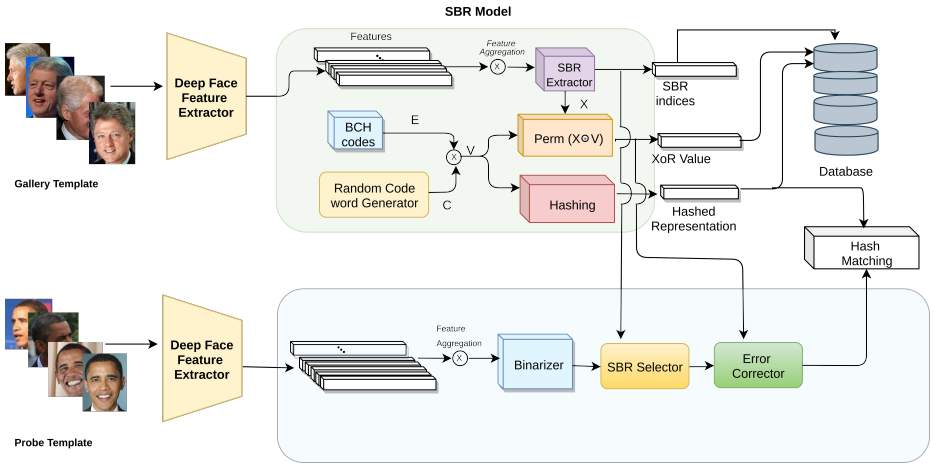

Figure 1. Illustration of proposed method. Significant Binary Representation (SBR) extractor is used to extract the significant features. SHA1 based hashing technique is applied on the unique code word generated for each gallery template

with the original biometric data and it has been shown that approaches similar to [11] can be used to recover user specific information. This necessitates a need for a representation that is disentangled from the original biometric data. This will help reduce the extent to which any information about the original biometric data can be recovered, given the new representation.

One approach to creating the aforementioned representation can be similar to [26] [15] where each user is assigned a maximum entropy binary (MEB) code. More precisely, let $c_i \sim B(1, 0.5)$ be the binary variable for each bit of the code, where $B(1, 0.5)$ is the maximum entropy Bernoulli distribution, and the resultant MEB code with independent bits is $MEB_i = [c_1, c_2, ..., c_N]$. The deep network learns the mapping function $Z : I \mapsto MEB_i$. where I is the input face image, MEB_i is the d dimensional binary representation assigned to the i^{th} user and N is the maximum number of users handled by the system. This representation provides the necessary disentanglement as it has no explicit correlation with the original biometric modality (the users face). Also the new representation provides the cancelability aspect to the biometric template.

The above-mentioned approach[26] that learns the mapping from input biometric modality to the MEB feature embedding necessitates allocation of these MEB features to all the users of the biometric system upfront. The discriminability of the approach is enforced by the assignment of the user specific maximum entropy bits (which by definition is maximally separated in the feature space) and the deep network merely unearths a transformation of the input face image satisfying the constraint. However, for many applications, this method may not be practical because of the incremental nature in which a real-world biometric system is scaled. Often such applications would require adding new users to the biometric system as the system scales. Handling dynamic enrollment of new users would require: a) possibly redefining N which might additionally entail increasing the MEB representation dimension d; and/or b) relearning of the mapping between input space and the MEB feature representation for the entire set of users, so as to accommodate the new identities while maintaining separability. Furthermore, it would be necessary to learn new MEB representations for the entire user base even if the stored MEB representation corresponding to only one of the users is compromised. Another limitation from a pragmatic of view of this representation is that, in order to incorporate the method to work with any modern face recognizer, it would take retraining the entire system from scratch. This would be not feasible for real world applications which are already deployed.

3.2. Significant Feature Representation

With these observations in mind and in order to create a robust representation, we analyze the feature embedding of a face recognizer $f_i \in R^d$. As mentioned earlier, two feature embeddings which are similar lie in close proximity in the d dimensional manifold. Let us consider three feature embeddings f_i, f_j ,f_k corresponding to images I,J,K. Assume images I and J belong to the same identity and image K belongs to a different identity. So, feature embedding f_i and f_j lie relatively close to each other on the manifold as compared to f_k. Mathematically, $CosSim(f_i, f_j) > CosSim(f_i, f_k)$ and $CosSim(f_i, f_j) > CosSim(f_j, f_k)$. Analyzing the feature embedding f_i which is unit normalized, each value in the feature embedding represents presence or absence of a direction in the d dimensional man-

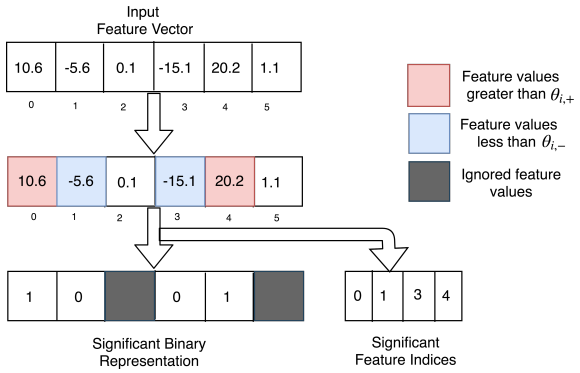

Figure 2. Illustration of the Significant Binary Representation (SBR) extractor used to extract the significant features.

ifold. A relatively large positive value indicates a strong presence of the feature, whereas a negative value indicates a strong absence of the same. The proximity of feature embedding in the manifold as given by cosine similarity indicates the correlation in the directions of feature values in the feature embedding. These decisive features are important in the feature embedding. Figure 1 illustrates the overall proposed architecture

The extraction of significant features in [18] and [29] relied on the difference between person specific and common database mean feature values normalized with respect to person specific sample standard deviation of feature values. Such calculations assume the availability of multiple feature vector samples during enrollment for each person. But this assumption is rather restrictive in real life applications, and in our database, for example, a large proportion of templates is constructed from a single feature vector.

Due to the proximity of genuine templates in the d dimensional manifold after training, the large positive or large negative feature values standing further away from 0, are likely to remain in the same area for genuine template pairs. But, since we are not able to calculate the sample standard deviation, we simply rely on two thresholds $\theta_{i,-}$ and $\theta_{i,+}$; if the jth feature value of template i, $f_{i,j}$, is less than $\theta_{i,-}$ or greater than $\theta_{i,+}$, then it is considered to be a significant feature.

We select the thresholds $\theta_{i,-}$ and $\theta_{i,+}$ for each template i such that a *constant* number of features ($X\%$ of the total number of features) in the template lesser and greater than these thresholds are chosen. This allows us to obtain binary vectors of the same dimension and having same number of 0 and 1 bits. Note that there is a trade-off in the number of selected significant features: on the one hand, we want a smaller number of most reliable bits in the resulting binary vector which will reduce the error correction load and increase the matching performance; and on the other hand, we need sufficient number of bits to prevent brute force attacks on the secured templates.

This method is used to create a *significant binary representation* (SBR) of all the gallery templates. Along with these binary representations, the locations (indices) of the significant feature values in the feature vector are also stored. Figure 2 illustrates the SBR extraction procedure. While processing a probe template, we rely on a single threshold θ_p, to binarize the probe template feature vector. All the values in the probe feature embedding greater than θ_p is set to 1 and less than θ_p is set to 0. At the time of verification, we compare only the significant feature values of the gallery to the corresponding positions of the probe (using the stored indices for the gallery). In order to find the similarity between a gallery and probe templates created using the new significant features, we can define a new metric based on Hamming distance (H_m).

$$S = 1 - \frac{H_m(F_g, F_p)}{L_{idx}}$$

where L_{idx} corresponds to the number of significant values in the gallery feature vector(F_g) and F_p is the probe feature vector. But storing the significant bit representation as is, in a traditional database, may result in some information leakage. In order to avoid any leakage of information and to provide an additional layer of security for the templates created using significant feature representation, we employ a fuzzy commitment scheme.

3.3. Fuzzy commitment scheme

In order to secure the extracted binary template, we utilize the fuzzy commitment scheme [16]. Suppose that $\boldsymbol{b}_i$ is our significant binary vector extracted from a gallery template, We allocate $\boldsymbol{v}_i$ a unique valid code word randomly generated for each template i. This code word is capable of correcting t errors using an error correction method. Then, our secured gallery template is the set $\{\boldsymbol{b}_i \oplus \boldsymbol{v}_i, H(\boldsymbol{v}_i)\}$, where H is some non-invertible cryptographic hash function. Given a probe binary vector $\boldsymbol{b}'_i$ with distance from $\boldsymbol{b}_i$ less than t in the significant vector positions, we can perform error correction on $(\boldsymbol{b}_i \oplus \boldsymbol{v}_i) \oplus \boldsymbol{b}'_i$ to recover $\boldsymbol{v}_i$ and check that we get same hash $H(\boldsymbol{v}_i)$. To prevent the attacks via record multiplicity, we can also employ the random permutation technique described in [17]. In this technique, the $\boldsymbol{b}_i$ are permuted using a randomly generated permutation $\boldsymbol{P}_i$ before the application of the fuzzy commitment scheme; the parameters of permutation $\boldsymbol{P}_i$ are kept as auxiliary data in the secured template.

4. Experiments

In this section, we describe the datasets, experimental setup, evaluation methodology and also discuss the results

obtained by employing the significant feature representation.

4.1. Datasets

We conducted our experiments on two benchmark datasets, primarily consisting of unconstrained face images.

IJB-A. IARPA Janus Benchmark A (IJB-A) dataset consists of 5712 images and 2805 videos of 500 subjects. The IJB-A evaluation protocol consists of verification (1:1 matching) over 10 splits. Each split contains around 11,748 pairs of templates (1,756 positive and 9,992 negative pairs) on average. Ten random splits of training and testing are provided as part of the protocol of IJB-A. We report 1:1 verification results on the protocol provided in the test splits.

IJB-C. The IARPA Janus BenchmarkC (IJB-C) is a larger dataset, being an extension to the IJB-A dataset, with 3541 subjects which has around 31,334 still images and 117543 frames. Similar to IJB-A, we report 1:1 verification results on the protocol provided.

4.2. Evaluation metric

We used cosine similarity as the evaluation measure when analyzing probe and gallery features directly extracted from the face recognizer. On the other hand, exact hash matching is used to evaluate the fuzzy commitment based method. The 1:1 verification results are evaluated using the ROC curve and the TAR (True Accept Rate) performance is reported for different FAR (False Accept Rate) values.

4.3. Face Recognizer

We trained a deep convolutional neural network based on the ResNeXt architecture[34]. The network was pre-initialized with weights of a model trained on the Imagenet-1K dataset and then trained on a refined version of the MSCeleb-1M dataset of which about $55K$ subjects and $1M$ images were used. The network was trained for 32 epochs using mini batch stochastic gradient descent. The learning rate was reduced by a factor of 10 on [6,12,18,24] epochs. The network once trained, was then finetuned on the UMD-Faces dataset. In order to improve the discriminative ability and the performance, hard negative mining was done using the triplet loss function. Once completely trained, 128-dimensional features of face images of IJB-A and IJB-C were extracted. The last layer of the network was linearly activated so that the features are spread on the 128 dimensional manifold. This results in the values of the extracted features ranging from $-\infty$ to ∞. Aggregated gallery and probe templates were created by averaging the features which constituted the template. In the first row of Table 1 and Table 2, we provide the TAR rate at different FAR on the IJB-A and IJB-C dataset respectively using the original, unsecured features.

4.4. Significant Features

In accordance with the method discussed in section 3.2, we create a binarized representation from the extracted features by analyzing the significant features within a feature vector. In order to find these significant features, we aggregate the gallery features by finding the mean of extracted facial image features present in a given template. Once the aggregated feature vector corresponding to a gallery is created, significant features are identified by sorting the feature vector and identifying the smallest and largest set of values in this feature vector. We empirically determined that keeping lowest and highest 25% of features as significant features gave a good trade-off between the error correction load and performance. We assign 0 to the lowest set of significant features and 1 to the highest set of significant features. We also record the positional information of these features by recording their indices in the feature vector. The rest of the insignificant features are ignored.

To avoid any leak in information we use a fuzzy commitment scheme to protect the feature template created using the significant features. We randomly create a N-bit long random code word for each gallery template. We follow the process explained in section 3.3 for comparing features. N is chosen so that the code's payload could accommodate our binary vector of consistent features (64 bits), and able to correct a specified number t of erroneous bits. For example, if we set our fuzzy commitment scheme to be able to correct 10-bit error, then we take BCH code of length 127 with 64 bits of payload data and 10 bit correction capability. We employ a SHA1 based hashing method to hash the corresponding cord word for each gallery template. Additional auxiliary information is also stored in the database. The probe template is binarized by sorting the probe feature vector and by setting the highest 50% values to 1 and remaining 50% to 0. The process of matching is done as explained in section 3.3. Similar to the original features in Table 1 and Table 2, the second row lists the TAR rate at different FAR for IJB-A and IJB-C datasets when using a hamming distance based metric on the significant binary representation. The third row list the TAR rate at various FAR after incorporating the fuzzy commitment scheme.

Table 1. IJB-A 1:1 Verification TAR(%)

Method	FAR			
	10^{-1}	10^{-2}	10^{-3}	10^{-4}
Original unsecured	97.96	95.83	92.64	85.65
SBR	97.49	94.05	90.27	83.92
SBR with hashing	97.64	94.31	89.77	84.04

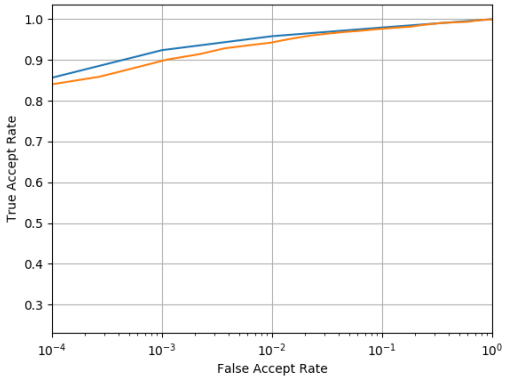

Figure 3. ROC curve for 1:1 Verification IJB-A (In log scale). Blue curve corresponds to the ROC of original, unsecured features. Orange curve corresponds to the ROC of the features secured using significant bit representation with fuzzy commitment. (Best viewed in color)

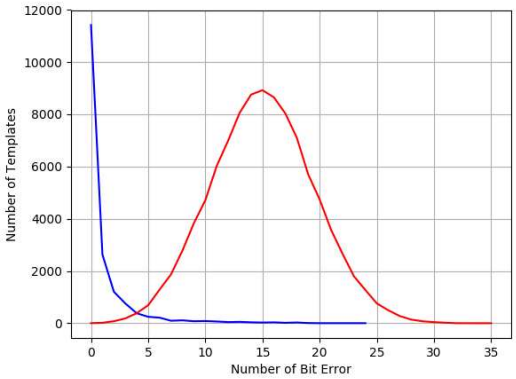

Figure 5. Bit error distribution analysis of genuine and imposter template pairs. Blue curve indicates the distribution of genuine template pairs and red curve indicates the distribution of imposter template pairs. Optimal separation can be seen at 4 bit error. (Best viewed in color)

Table 2. IJB-C 1:1 Verification TAR(%)

Method	FAR				
	10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}
Original unsecured	98.37	96.02	92.40	86.89	74.79
SBR	97.54	94.49	89.98	81.00	66.36
SBR with hashing	97.88	95.04	90.01	81.00	62.75

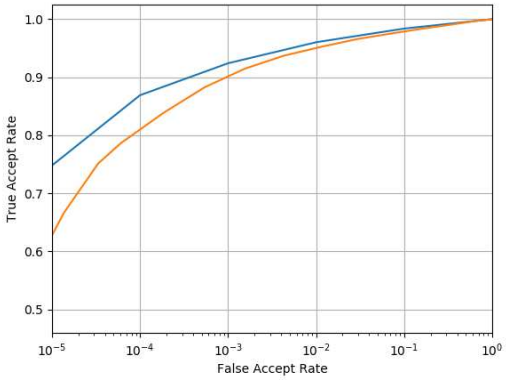

Figure 4. ROC curve for 1:1 Verification IJB-C (In log scale). Blue curve corresponds to the ROC of original, unsecured features. Orange curve corresponds to the ROC of the features secured using significant bit representation with fuzzy commitment. (Best viewed in color)

5. Discussion

The results of our experiments show the practicality and ease of incorporating privacy preserving significant features and fuzzy commitment scheme to any ordinary face recognizer. This helps improve the robustness of using the face as a biometric without compromising private information in case of a stolen template. We acknowledge that there is a performance drop compared to the original unsecured features. One of the reasons for this may be, that when significant features are created, the degree of freedom of each value ($-\infty$ to ∞) is restricted to just two values $\{0, 1\}$. Another reason for information loss and performance reduction can be attributed to the reduction in feature size when creating the significant features. Increasing the feature dimension of the face recognizer, not only has the potential to increase the performance but also could provide better security against brute force attacks. The decision to consider a face recognizer having 128 dimensional feature vector for the experiments was primarily due to the observation that most of existing state-of-the-art face recognizers had a similar feature dimension. Moreover, one of the primary objectives of the paper was to explore an approach to incorporate security into existing face recognizers, rather than advocating an alternative method to design one.

Figure 5 shows the difference in the number of bit mismatches while comparing genuine template pairs and imposter template pairs in the IJB-A dataset. For this analysis, all the templates were binarized by identifying the lowest and highest 25% feature values and comparing the binary values in the corresponding positions in the templates. The number of bit difference between a genuine template pair is substantially less than imposter template pair. This helps in restricting the error correction capability of the fuzzy commitment scheme to a lower number, so that only the genuine template pairs get matched.

Another interesting fact to note is that, the dataset used

in [25] [26] are small datasets with less number of identities and limited pose variance. We report the results of the experiment on two unconstrained image datasets. The helps reinforce the fact that the proposed method can be easily adapted into existing state-of-the-art face recognizers. Additionally our method doesn't require end-to-end retraining of the face recognizer, making it a plug-and-play framework.

Furthermore, integration of fuzzy commitment scheme provides the cancelability aspect to the biometric templates. Any compromise in the security of database, resulting in a stolen template, can be addressed by merely changing the random code word associated with each template.

6. Conclusion and Future Work

We present a novel approach for incorporating template level security into an existing face recognizer. By adopting fuzzy commitment scheme, we integrate security guarantees of the scheme into the face recognizer without significantly affecting its performance. The design of the method allows it to be easily integrated into any existing face recognizer. In order to underline the scalability and broader applicability of our approach, we report results on two unconstrained face datasets with large number of identities.

One possible way to improve the presented approach will involve training of a face recognizer best suited for significant feature representation. For example, in addition to the currently used triplet loss optimization, one could use loss terms penalizing for significant feature instability or flipping, bias in significant feature distributions, etc. Ideally, the network training would account for the particular way of computing the scores in the final system with significant feature binarization.

7. Acknowledgement

This material is based upon work partially supported by the National Science Foundation under Grant IIP #1822190

References

[1] T. Ahonen, A. Hadid, and M. Pietikainen. Face description with local binary patterns: Application to face recognition. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (12):2037–2041, 2006. 2

[2] M. Ao and S. Z. Li. Near infrared face based biometric key binding. In *International Conference on Biometrics*, pages 376–385. Springer, 2009. 2

[3] A. Bansal, A. Nanduri, C. D. Castillo, R. Ranjan, and R. Chellappa. Umdfaces: An annotated face dataset for training deep networks. In *Biometrics (IJCB), 2017 IEEE International Joint Conference on*, pages 464–473. IEEE, 2017. 2

[4] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. In *European Conference on Computer Vision*, pages 43–58. Springer, 1996. 2

[5] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman. Vggface2: A dataset for recognising faces across pose and age. In *Automatic Face & Gesture Recognition (FG 2018), 2018 13th IEEE International Conference on*, pages 67–74. IEEE, 2018. 2

[6] K.-Y. Chee, Z. Jin, D. Cai, M. Li, W.-S. Yap, Y.-L. Lai, and B.-M. Goi. Cancellable speech template via random binary orthogonal matrices projection hashing. *Pattern Recognition*, 76:273–287, 2018. 2

[7] B. Chen and V. Chandran. Biometric based cryptographic key generation from faces. In *Digital Image Computing Techniques and Applications, 9th Biennial Conference of the Australian Pattern Recognition Society on*, pages 394–401. IEEE, 2007. 2

[8] Y. Chen, Y. Wo, R. Xie, C. Wu, and G. Han. Deep secure quantization: On secure biometric hashing against similarity-based attacks. *Signal Processing*, 154:314–323, 2019. 2

[9] J. Deng, J. Guo, and S. Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *arXiv preprint arXiv:1801.07698*, 2018. 2

[10] Y. C. Feng and P. C. Yuen. Binary discriminant analysis for generating binary face template. *IEEE Transactions on Information Forensics and Security*, 7(2):613–624, 2012. 2

[11] M. Fredrikson, S. Jha, and T. Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1322–1333. ACM, 2015. 1, 3

[12] Y. Guo, L. Zhang, Y. Hu, X. He, and J. Gao. Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. In *European Conference on Computer Vision*, pages 87–102. Springer, 2016. 2

[13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2

[14] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 2

[15] A. K. Jindal, S. Chalamala, and S. K. Jami. Face template protection using deep convolutional neural network. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 575–5758. IEEE, 2018. 2, 3

[16] A. Juels and M. Wattenberg. A fuzzy commitment scheme. In *ACM Conference on Computer and Communications Security*, pages 28–36, 1999. 4

[17] 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. *Information Forensics and Security, IEEE Transactions on*, 6(1):107–121, 2011. 4

[18] T. A. Kevenaar, G. J. Schrijen, M. van der Veen, A. H. Akkermans, and F. Zuo. Face recognition with renewable and privacy preserving binary templates. In *Automatic Identification Advanced Technologies, 2005. Fourth IEEE Workshop on*, pages 21–26. IEEE, 2005. 1, 2, 4

[19] Y. Kim and K.-A. Toh. A method to enhance face biometric security. In *Biometrics: Theory, Applications, and Systems, 2007. BTAS 2007. First IEEE International Conference on*, pages 1–6. IEEE, 2007. 2

[20] C. Liu and H. Wechsler. Gabor feature based classification using the enhanced fisher linear discriminant model for face recognition. *IEEE Transactions on Image processing*, 11(4):467–476, 2002. 2

[21] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphereface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017. 2

[22] H. Lu, K. Martin, F. Bui, K. Plataniotis, and D. Hatzinakos. Face recognition with biometric encryption for privacy-enhancing self-exclusion. In *Digital Signal Processing, 2009 16th International Conference on*, pages 1–8. IEEE, 2009. 1

[23] J. Lu, Y.-P. Tan, and G. Wang. Discriminative multimanifold analysis for face recognition from a single training sample per person. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):39–51, 2013. 2

[24] D. C. Ngo, A. B. Teoh, and A. Goh. Biometric hash: high-confidence face recognition. *IEEE transactions on circuits and systems for video technology*, 16(6):771–775, 2006. 2

[25] R. K. Pandey and V. Govindaraju. Secure face template generation via local region hashing. In *Biometrics (ICB), 2015 International Conference on*, pages 299–304. IEEE, 2015. 1, 2, 7

[26] R. K. Pandey, Y. Zhou, B. U. Kota, and V. Govindaraju. Deep secure encoding for face template protection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 77–83. IEEE, 2016. 1, 2, 3, 7

[27] R. Ranjan, C. D. Castillo, and R. Chellappa. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507*, 2017. 2

[28] N. K. Ratha, S. Chikkerur, J. H. Connell, and R. M. Bolle. Generating cancelable fingerprint templates. *IEEE Transactions on pattern analysis and machine intelligence*, 29(4):561–572, 2007. 2

[29] V. Talreja, M. C. Valenti, and N. M. Nasrabadi. Multibiometric secure system based on deep learning. In *Signal and Information Processing (GlobalSIP), 2017 IEEE Global Conference on*, pages 298–302. IEEE, 2017. 2, 4

[30] A. B. Teoh, D. C. Ngo, and A. Goh. Personalised cryptographic key generation based on facehashing. *Computers & Security*, 23(7):606–614, 2004. 2

[31] A. B. J. Teoh and C. T. Yuang. Cancelable biometrics realization with multispace random projections. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 37(5):1096–1106, 2007. 2

[32] M. Van Der Veen, T. Kevenaar, G.-J. Schrijen, T. H. Akkermans, and F. Zuo. Face biometrics with renewable templates. In *Security, Steganography, and Watermarking of Multimedia Contents VIII*, volume 6072, page 60720J. International Society for Optics and Photonics, 2006. 1

[33] Y. Wu and B. Qiu. Transforming a pattern identifier into biometric key generators. In *2010 IEEE International Conference on Multimedia and Expo*, pages 78–82. IEEE, 2010. 2

[34] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 5

[35] Y. Zheng, D. K. Pal, and M. Savvides. Ring loss: Convex feature normalization for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5089–5097, 2018. 2