Secure Face Matching Using Fully Homomorphic Encryption

Vishnu Naresh Boddeti
Michigan State University, East Lansing MI 48824
vishnu@msu.edu

Introduction: A typical face recognition system acquires a facial image, which undergoes some pre-processing, following which high-dimensional features are extracted. During enrollment, these features vectors are stored in a database along with their identity labels. This database is then used to verify a person's claimed identity (face verification) or determine a person's identity (face identification).

The database of face templates is a prime candidate for malicious attacks. Directly storing the facial feature vectors in the database can significantly compromise the privacy of the enrolled users and the security of the authentication system. For instance, it has been demonstrated recently that access to the representations can enable the reconstruction of a user's original facial image [4] by a malicious attacker. The reconstructed face images were shown to be of sufficient quality to be successfully matched by a state-of-the-art face recognition system. Similarly, soft facial attributes such as age, gender, ethnicity etc. can be predicted with high accuracy from facial features [3]. Therefore, it is imperative to devise techniques to prevent information leakage from face representations while preserving face matching performance, thereby preventing attacks of this nature.

An attractive solution for secure face matching is the use of *homomorphic cryptosystems* to protect the face template database and perform matching over encrypted data. Pursuantly, we present a *fully homomorphic encryption* based approach to (1) cryptographically secure the database of face templates and the feature vector of the presented face image, and (2) perform matching directly in the encrypted domain without the need for decrypting the templates. We leverage the observation that a typical face matching metric, either Euclidean distance or cosine similarity, can be decomposed into its constituent series of addition and multiplication operations. Both of these operations are supported by FHE schemes over encrypted data.

The key technical barrier to realizing homomorphic encryption based face matching is the computational complexity of homomorphic encryption, especially homomorphic multiplication. A straight-forward application of FHE for face matching needs 48.7 MB of memory for each 512-dimensional encrypted face template and 12.8 secs for matching a single pair of such templates. To address this limitation we present a scheme: (1) Utilize a more efficient FHE scheme, namely Fan-Vercauteren [1], easing the computational burden to 16.5 MB of memory and 0.6 secs per pair of templates. These requirements could still be prohibitive for practical deployment. (2) Utilize a batching scheme that allows homomorphic multiplication of multiple values at the cost of a single homomorphic multiplication. This scheme reduces the computational requirements to, 16 KB memory per template and 0.01 secs to match a pair of templates. (3) Dimensionality reduction to further provide a trade-off between computational efficiency and matching performance.

Main Contribution: The score between two d-dimensional representations (either Euclidean distance or cosine dissimilarity) is first decomposed into its constituent arithmetic operations, d additions and d multiplications. While FHE schemes can be directly utilized for computing the score in the encrypted domain, it suffers from high computational requirements. To overcome this, we utilize a batching scheme for encoding an entire vector of real values instead of each of its elements. Such and encoding allows us to amortize multiple homomorphic multiplications within a single homomorphic multiplication, thereby providing significant computational efficiencies.

Experimental Results: For our experimentation, we consider two different state-of-the-art face representations, namely, FaceNet [5] and SphereFace [2] with 128 and 512 dimensional features, respectively. We evaluate these representations over multiple datasets of varying complexity, namely, LFW and IJB-B datasets.

For a given desired level of security, Table 1 shows the effect of the different parameter settings and the dimensionality of the facial features on the computational complexity of face matching over encrypted features. We report a breakdown of the time taken for each of the steps in the matching process, the total matching time and the size of the templates. We observe that the fastest FHE based face matching is admittedly over 1000x slower than matching in the real-domain. Nevertheless, batching can offer a significant speed-up and memory savings over element-

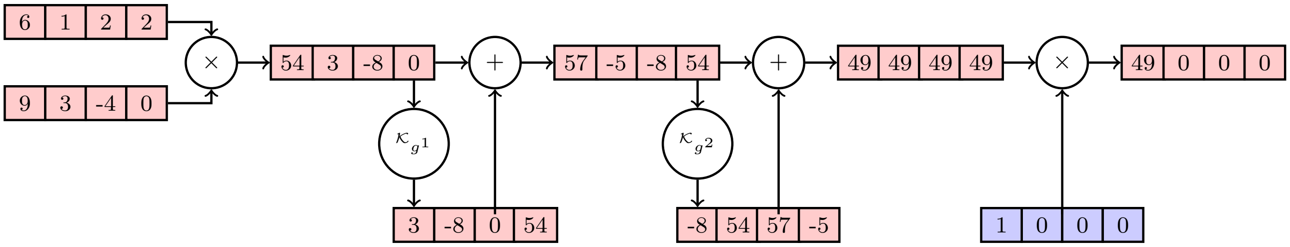

Figure 1: Homomorphic computation of inner product between vectors using batching. In this process multiple elements are encrypted into a single encrypted block, preventing access to each element of the encrypted block. After the Hadamard product, the sum of the elements within the encrypted block can be computed through repeated cyclic rotations and additions.

Table 1: Security Parameters, Timing and Memory

Security	Dim	No FHE		Parameters			Batching					No Batching				
in bits		Time	Mem				Time (ms)				Mem	Time (ms)				Mem
(λ)	(d)	(μs)	(KB)	n	t (bits)	q	Enc	Score	Dec	Total	(KB)	Enc	Score	Dec	Total	(MB)
128	64	0.44	2.0	128	110	40961	0.07	0.17	0.01	0.25	2.0	4.40	5.25	0.01	9.66	0.25
	128	0.89	4.0	256	110	40961	0.14	0.38	0.02	0.59	4.0	17.57	21.05	0.02	38.64	1.0
	512	3.48	16.0	1024	110	40961	0.58	1.80	0.07	2.45	16.0	280.19	343.81	0.08	624.07	16.5
	1024	7.49	32.0	2048	110	40961	1.14	4.02	0.15	5.80	33.0	1135.44	1411.82	0.16	2547.42	66.0
	1024	7.49	32.0	4096	110	40961	2.27	8.36	0.30	11.42	66.0	2214.88	2924.75	0.33	5139.97	131.0

Table 2: Face Recognition Accuracy (TAR @ FAR in %)

Dataset	Method	128-D FaceNet			512-D SphereFace			64-D PCA FaceNet			64-D PCA SphereFace		
		0.01%	0.1%	1%	0.01%	0.1%	1%	0.01%	0.1%	1%	0.01%	0.1%	1%
LFW	No FHE	84.06	94.56	98.65	90.49	96.74	99.11	83.99	94.64	98.73	88.41	95.80	98.87
	FHE	84.05	94.56	98.65	90.49	96.74	99.11	84.00	94.64	98.72	88.38	95.80	98.87
IJB-B	No FHE	25.77	48.31	74.47	7.86	31.27	69.83	24.95	47.80	74.58	7.13	29.72	69.79
	FHE	25.78	48.28	74.46	7.86	31.27	69.82	25.06	47.78	74.66	6.85	29.70	69.69

wise homomorphic face matching. In comparison, direct application of FHE for face matching requires ∼12.8 secs per match pair and 48.7 MB for each template for 512-dimensional features. In contrast, our results, suggest that face matching over encrypted templates using the proposed FHE framework can provide high levels of security, 128 bits, and real-time matching over a small database of face templates.

We report the results of face matching experiments on benchmark datasets for FaceNet [5] and SphereFace [2] in Table 2. We report true acceptance rate (TAR) at three different operating points of 0.01%, 0.1% and 1.0% false accept rates (FARs). We first evaluate face matching performance over unencrypted face templates as a baseline to compare against. Results indicate that homomorphic face matching over encrypted features can perform as well as matching raw features while providing template protection, preventing information leakage and preserving the privacy of the users. Finally, we observe that even simple PCA based dimensionality reduction can perform comparably to the original high-dimensional features, while affording significant homomorphic face matching efficiency.

References

[1] J. Fan and F. Vercauteren. Somewhat practical fully homomorphic encryption. *IACR Cryptology ePrint Archive*, 2012:144, 2012.

[2] W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*, volume 1, 2017.

[3] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, 2015.

[4] G. Mai, K. Cao, P. C. Yuen, and A. K. Jain. Face image reconstruction from deep templates. *arXiv preprint arXiv:1703.00832*, 2017.

[5] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, pages 815–823, 2015.