

On the Protection of Face Recognition Embeddings

Gábor György Gulyás

Vitarex Stúdió Ltd

Budapest, Hungary

gabor@gulyas.info

Abstract—Face recognition technologies rely on face embeddings, numerical representations of biometric features, which are increasingly used in security and commercial applications. However, the privacy risks associated with storing and processing these embeddings are significant, particularly under strict regulations such as the GDPR and the AI Act. This paper investigates two major techniques for protecting face embeddings: Locality-Sensitive Hashing (LSH) and Homomorphic Encryption (HE). Through a case study using random projections and the LFW dataset, we show that while LSH allows great reduction in data sizes and offers efficient approximate matching, it provides weak resistance to re-identification attacks. In contrast, HE enables computation directly on encrypted data and offers a more secure, though computationally expensive, alternative. We evaluate recent HE-based approaches and propose optimizations.

Keywords—face recognition, biometric templates, privacy, security.

I. INTRODUCTION

Face recognition technologies are rapidly becoming integral components of modern smart city infrastructures and daily digital interactions. From secure access control in buildings and devices to personalized services in businesses, the ability to reliably identify individuals through facial biometrics offers both convenience and efficiency. At the core of these systems lie face embeddings — mathematical representations of facial features—serving as the fundamental units for recognition of individuals.

As their use becomes more widespread, security and privacy of face embeddings emerge as critical concerns. Unlike passwords or tokens, biometric data is inherently sensitive and immutable: once compromised, it cannot be revoked or changed. Hence there are strict provisions under both the GDPR [1] and the AI Act [2]. Protecting face embeddings from exposure, inversion, and misuse is therefore essential for maintaining user trust and ensuring the long-term viability of face recognition systems.

What are face embeddings? Briefly, they are our facial fingerprints; embeddings belonging to a specific person uniquely describe that individual, but differ notably from the embeddings of other individuals. State-of-the-art face recognition systems are predominantly based on deep learning, particularly convolutional neural networks (CNNs). These models learn to extract compact and discriminative feature vectors (known as face embeddings) that capture the unique identity of individuals while remaining invariant to variations in pose, lighting, and expression. Face recognition technologies use various deep learning architectures, and their vector structure usually consists




Figure 1. Individuals' embeddings (with at least 10 embeddings per individual) visualized from the Labeled Faces in the Wild dataset [3].

of float values with a variable length, having typically 128-512 dimensions.

We provide an example on the behavior of embeddings in fig. 1, where we visualize individuals' embeddings from the Labeled Faces in the Wild dataset [3], using TSNE [4]. It is clearly visible that each individual's embedding cluster together but separately from others.

This paper investigates how state-of-the-art technical security measures can protect face embeddings [5]; in particular, we analyze locality-sensitive hashing (LSH) and homomorphic encryption (HE). LSH offers transformation with information loss to map embeddings into distance-preserving hashes that can still support approximate matching. Homomorphic encryption, on the other hand, provides strong cryptographic guarantees by enabling computation on encrypted data without ever revealing the plaintext embeddings.

We examine the effectiveness of LSH through an experiment on random projection using the LFW dataset. Our results show that while LSH can preserve identity similarity for efficient matching, it offers limited resistance to re-identification under adversarial conditions. To address these shortcomings, we investigate homomorphic encryption as a cryptographic alternative that allows secure computation on encrypted embeddings.

The rest of this paper is organized as follows. In Section II, we introduce Locality-Sensitive Hashing and evaluate its security and utility in the context of face embeddings. Section

III presents Homomorphic Encryption, reviewing relevant literature and comparing the effectiveness of recent approaches, also concerning face recognition. Finally, Section IV concludes the paper with a summary of our findings and outlines directions for future research.

II. LOCALITY-SENSITIVE HASHING

Locality-Sensitive Hashing (LSH) is a technique used to reduce data dimensionality while retaining the ability to efficiently perform approximate nearest neighbor searches. Unlike traditional cryptographic hash functions, which aim to maximize randomness and unpredictability, LSH functions are designed so that similar inputs (e.g. according to cosine distance) are more likely to be hashed to nearby buckets. This property makes LSH particularly useful for tasks like face recognition, where exact matches are rare, but similar embeddings need to be quickly identified.

A. Related Work

Works on face hashing derive their own features for creating embeddings. One of the first works on hashing faces introduces Local Feature Hashing (LFH), a face recognition method designed for scalability and robustness in large-scale, real-world applications [6]. LFH combines local image descriptors with p-stable distribution-based Locality Sensitive Hashing (pLSH) to achieve fast and accurate approximate nearest neighbor (ANN) searches.

Another, more recent work, explores the effectiveness of various Locality-Sensitive Hashing (LSH) techniques for improving the performance compared to previous LFH methods [7]. Due to the increasing volume of facial images in social networks and databases, reducing system response time is critical. The authors propose an unsupervised face recognition pipeline using local feature extraction, PCA for dimensionality reduction, and several LSH variants for efficient indexing and similarity search. Their method is evaluated on three facial image datasets, and hashing function evaluated concerning their impact, scalability and responsiveness. Their work suggest future research into combining LSH with deep learning.

B. Deep Learning and LSH

Since embeddings are structured simply as float vectors, random projection [8] seems a suitable LSH technique for dimension reduction. Therefore, random projection is a technique for reducing the dimensionality of high-dimensional data while approximately preserving the pairwise distances between points. Its core idea is based on the Johnson–Lindenstrauss lemma, which states that a small set of points in high-dimensional space can be projected into a lower-dimensional space in such a way that the distances between the points are nearly preserved.

Mathematically, given a data vector $x \in R^d$, we multiply it by a random matrix $R \in R^{k \times d}$, where $k \ll d$, to obtain a lower-dimensional representation $x' = Rx$. The entries of R are typically sampled from a Gaussian distribution $\mathcal{N}(0, 1)$ or a sparse distribution with similar properties. This method is computationally efficient and preserves the structure of the data well enough for many applications.




Figure 2. Our methodology: each profile (rows) were split half, then we used different LSH hashes to map their embeddings from 2048 bits to 64 bits.

C. Evaluation

In this section we evaluate random projection for face recognition embeddings. Our methodology is as follows. By using a subset of the Labeled Faces in the Wild (LFW) dataset [3], we evaluate if identity-preserving similarity can be maintained after applying random projections (with different key matrices) that significantly reduce the dimensionality of the embeddings. One could imagine, that using a method with such information loss (e.g. shrinking embeddings only to their $1/32$), would prevent matching identities.

We first selected a subset of the LFW dataset by filtering for individuals with at least 40 available face embeddings, then we randomly sampled exactly 40 embeddings, ensuring balanced coverage across individuals. These embeddings were then split into two disjoint sets of 20, simulating two independent profiles of each person. This resulted in a dataset of 19 people.

To simulate an LSH scenario using random projection, we generated two independent random matrices $R_A \in R^{4 \times 128}$ and $R_B \in R^{4 \times 128}$, with entries sampled from a standard Gaussian distribution $\mathcal{N}(0, 1)$. These matrices were used to map each set of embeddings to a 64-bit binary hash, effectively reducing the 128-float (i.e., 2048 bits) representation to a compact binary code. We denote these LSH_1 and LSH_2 , and the resulting hashed profiles are set A and B, respectively (see in fig. 2).

This setup allows us to analyze the linkability properties of randomly projected embeddings across two independent LSH spaces (assuming not knowing the projection matrices). Furthermore, this evaluation models practical scenarios where an attacker may try to correlate different embedding databases. Hence, the underlying questions is: is using LSH secure under the condition that an attacker cannot access the projection matrix?

The baseline approach could be to guess the mapping between embeddings $e_1 \in |A|$, $e_2 \in |B|$. The probability of a correct mapping is $1/19 \sim 5.2\%$. A more sophisticated attack would be to map an embedding to the person to which it is more similar using cosine similarity: this naive algorithm also takes us to a re-identification rate of 5.2%.

The underlying problem is identical to a re-identification situation. The core component of re-identification algorithms is the similarity metric that determines if two pieces of information (e.g. profile descriptors) are related or not. Only




Figure 3. Hashed subset of the LFW dataset with only 20 embeddings per individual. Despite the significant loss of information, the hashing preserved differences between identities, while embeddings belonging to the same person are still clustered.

changing the similarity metric without touching the rest of the algorithm can yield significantly better results [9].

In our next step, we conducted a classification experiment using XGBoost, a gradient-boosted decision tree model. We constructed a binary classification task where the input consisted of embedding pairs, and the labels indicated whether the pair belonged to the same identity (positive class) or to different identities (negative class). The training and test data structure is illustrated in fig. 2.

Due to the inherent class imbalance in pairwise face verification tasks, where negative pairs typically outnumber positive ones (e.g. 4104 vs 228 in our experiment), we applied class-balancing techniques during training. Specifically, we assigned higher importance to underrepresented classes and reinforced the model’s sensitivity to the minority class. The model was trained on the train set of embedding pairs and evaluated on a disjoint test set (with a 80 : 20 train-test ratio). With this setup, we have achieved an overall 89.87% accuracy, with a recall of 91% for non-matching hashed embedding pairs (4104 cases), and 68% for matching ones (228 cases).

This leads us to the conclusion that while Locality-Sensitive Hashing is an efficient tool for reducing data volume while preserving utility, but alone itself offers weak utility.

III. HOMOMORPHIC ENCRYPTION

The previous conclusion on LSH leads us to the need to look for cryptographic solutions. This brings us to investigate Homomorphic Encryption (HE), which is an emerging cryptographic technique that enables computation on encrypted data without requiring decryption. In the context of face recognition, this means that sensitive biometric embeddings can remain encrypted while still supporting operations such as similarity comparison. This capability is especially valuable

for privacy-preserving applications, where biometric data must be processed by untrusted systems or outsourced to third-party cloud services - or simply just for compliance reasons.

Formally, a homomorphic encryption scheme allows specific algebraic operations (e.g., addition or multiplication) to be performed on ciphertexts such that, when decrypted, the result matches the operation as if it had been performed directly on the plaintexts. Fully Homomorphic Encryption (FHE) schemes support arbitrary computations on encrypted data, while Partial Homomorphic Encryption (PHE) schemes restrict the number or type of operations. The first working FHE solution was published in 2009 [10].

Although HE offers strong privacy guarantees, their use in face recognition systems has been limited due to computational overhead and complexity. However, recent advancements in the efficiency and usability of HE libraries have opened new possibilities for secure biometric matching, making it a promising direction for embedding protection in real-world deployments. In the following we provide a literature survey on new results.

Mi et al. [11] propose a privacy-preserving face recognition method that transforms facial data extracted from images in a way that conceals the subject’s identity, while still enabling accurate recognition. Their approach demonstrates strong recognition performance and improved privacy compared to earlier transformation-based techniques. However, one notable limitation is that, in the event of a data breach, the transformed representations can still be partially inverted—allowing the recovery of visual cues such as skin tone or hair color. Leaking such personal attributes can still lead to re-identification as demonstrated in [12], therefore any meaningful reconstruction or inference by an attacker must be prevented.

A recent study by Serengil et al. [13] presents an encryption-based face recognition approach in which most of the computation—including face detection and matching—is carried out directly in the encrypted domain. This design ensures a high level of security by minimizing exposure of sensitive data. The authors report that the system achieves both high accuracy and practical performance. While their method is sound, significant improvement needs to be done to reduce runtime efficiency. We have checked their approach by running our own simulations (with a different HE library), and we believe speedups of 4-8x could be possible, depending on the configuration – which we leave for future work.

A similar technique had been introduced before by Boddeti [14], who applies similar cryptographic optimizations to enable secure face recognition. However, their method has only been evaluated on isolated datasets in experimental settings. As a result, its applicability in real-world scenarios remains uncertain.

In summary, homomorphic encryption represents a compelling solution for protecting face recognition embeddings, offering strong theoretical guarantees and the ability to perform computations on encrypted face recognition embeddings. While prior work has laid important foundations, many of the proposed systems either fall short in terms of efficiency or

have yet to be validated in realistic, large-scale scenarios.

IV. CONCLUSION AND FUTURE WORK

As face recognition technologies become increasingly embedded in everyday digital infrastructure, the need for robust protection of biometric data, and specifically face embeddings, grows ever more urgent. In this paper, we examined the security and privacy challenges associated with storing and processing face recognition embeddings, with a particular focus on two technical approaches: Locality-Sensitive Hashing (LSH) and Homomorphic Encryption (HE).

Our evaluation of LSH, including a random projection-based case study, revealed that while LSH can efficiently reduce data size (up to 1/32) and preserve similarity for approximate matching, it provides no privacy protection against re-identification attacks. In contrast, Homomorphic Encryption offers a higher level of security by enabling computations to be carried out on encrypted embeddings. While existing HE-based face recognition systems demonstrate promising results in terms of privacy and accuracy, many still need future work on the practical side.

Future work should build on existing results by doing extended evaluations to larger and more life-like and diverse datasets, explore integration with secure hardware enclaves, and investigate hybrid schemes that combine LSH and encryption for a balanced trade-off between efficiency and security. Additionally, adapting these techniques to real-time video analytics remains a promising direction for achieving scalable, redundant, privacy-preserving biometric recognition for smart cities and enterprise applications.

V. ACKNOWLEDGEMENT

The author is grateful for Csaba Kiss, Dániel Kuknyó, Gergely Erdődi for their support provided while carrying out this work.

Project entitled "GDPR compliant anonymous facial recognition for analytics and improving security", with no. 2021-1.1.4-GYORSÍTÓSAV-2022-00076 has been implemented with the support provided by the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation Fund, financed under the 2021-1.1.4-GYORSÍTÓSAV funding scheme.

REFERENCES

- [1] *Regulation (eu) 2016/679 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/ec (general data protection regulation)*, <https://eur-lex.europa.eu/eli/reg/2016/679/oj>, Accessed: 2025-04-11, 2016.
- [2] *Regulation (eu) 2024/xxxx of the european parliament and of the council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain union legislative acts*, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>, Accessed: 2025-04-11, 2024.
- [3] G. B. Huang, M. Mattar, T. Berg, and E. Learned-Miller, "Labeled faces in the wild: A database for studying face recognition in unconstrained environments", in *Workshop on faces in 'Real-Life' Images: detection, alignment, and recognition*, 2008.
- [4] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne.", *Journal of machine learning research*, vol. 9, no. 11, 2008.
- [5] S. M. S. Ahmad, B. M. Ali, and W. A. W. Adnan, "Technical issues and challenges of biometric applications as access control tools of information security", *International journal of innovative computing, information and control*, vol. 8, no. 11, pp. 7983–7999, 2012.
- [6] Z. Zeng, T. Fang, S. Shah, and I. A. Kakadiaris, "Local feature hashing for face recognition", in *2009 IEEE 3rd International Conference on Biometrics: Theory, Applications, and Systems*, IEEE, 2009, pp. 1–8.
- [7] M. Dehghani, A. Moeini, and A. Kamandi, "Experimental evaluation of local sensitive hashing functions for face recognition", in *2019 5th International Conference on Web Research (ICWR)*, IEEE, 2019, pp. 184–195.
- [8] E. Bingham and H. Mannila, "Random projection in dimensionality reduction: Applications to image and text data", in *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining*, 2001, pp. 245–250.
- [9] G. G. Gulyás, B. Simon, and S. Imre, "An efficient and robust social network de-anonymization attack", in *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society*, 2016, pp. 1–11.
- [10] C. Gentry, "Fully homomorphic encryption using ideal lattices", in *Proceedings of the forty-first annual ACM symposium on Theory of computing*, 2009, pp. 169–178.
- [11] Y. Mi et al., "Privacy-preserving face recognition using trainable feature subtraction", in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 297–307.
- [12] I. Fábián and G. G. Gulyás, "De-anonymizing facial recognition embeddings", *Infocommunications Journal*, vol. 12, no. 2, pp. 50–56, 2020.
- [13] S. Serengil and A. Ozpinar, "Cipherface: A fully homomorphic encryption-driven framework for secure cloud-based facial recognition", *arXiv preprint arXiv:2502.18514*, 2025.
- [14] V. N. Boddeti, "Secure face matching using fully homomorphic encryption", in *2018 IEEE 9th international conference on biometrics theory, applications and systems (BTAS)*, IEEE, 2018, pp. 1–10.