

Enhancing Secure Cross-Border Collaboration among Law Enforcement Agencies for Facial Biometric Search

Kyriaki Miniadou

Institute of Computer Science (ICS)

Foundation for Research and Technology - Hellas (FORTH)

Heraklion, Greece

kminiadou@ics.forth.gr

Asterios Leonidis

Institute of Computer Science (ICS)

Foundation for Research and Technology - Hellas (FORTH)

Heraklion, Greece

leonidis@ics.forth.gr

Georgios Th. Papadopoulos

Institute of Computer Science (ICS)

Foundation for Research and Technology - Hellas (FORTH)

Heraklion, Greece

and

Department of Informatics and Telematics

Harokopio University of Athens

Athens, Greece

g.th.papadopoulos@hua.gr

Constantine Stephanidis

Institute of Computer Science (ICS)

Foundation for Research and Technology - Hellas (FORTH)

Heraklion, Greece

and

Department of Computer Science

University of Crete

Heraklion, Greece

cs@ics.forth.gr

Abstract—Addressing the growing challenge of combating international crime involves the development of secure and efficient methodologies that allow Law Enforcement Agencies to exchange information seamlessly, without being hindered by time-consuming bureaucratic processes. In this context, we present a solution centered on facial biometric search methodologies. Our approach underscores the importance of employing accurate and reliable methods to assess image data similarity, particularly in the domain of facial images, which pose unique challenges due to subtle variations. We propose a comprehensive solution that harnesses hashing techniques and homomorphic encryption. By doing so, our approach ensures secure data exchange while safeguarding confidentiality and integrity. We firmly believe that our approach will substantially improve collaboration in law enforcement efforts and make significant contributions to global security.

Index Terms—Biometrics, Deep Learning, Security, Interaction Issues, User Experience

I. INTRODUCTION

In the effort to enhance global security and combat criminal activities, Law Enforcement Agencies (LEAs) are increasingly collaborating and exchanging information across varied jurisdictions. This has necessitated the development of secure and robust mechanisms for data sharing while adhering to diverse legal frameworks and data privacy regulations.

To facilitate suspect identification and foster connections with ongoing investigations conducted by different international LEAs, we propose an approach to enhance collaboration centered on facial biometric search methodologies. This endeavor underscores the need of employing methodologies for

assessing image data similarity and facilitating the retrieval of relevant images with accuracy and reliability.

Our primary focus rests on a distinct facet of image similarity: that of facial images. Facial image analysis presents a unique and intricate challenge due to the subtle variations that exist within images of the same face. These variations encompass nuanced differences in pose, age, hair color, and even the presence or absence of facial hair.

The process of image similarity, confronts the daunting challenge of assessing similarity in high-dimensional datasets based on a query image. The complexity arising from the scalability as well as the high dimensionality of image data, poses a significant challenge in various applications that require such input. Addressing this challenge without compromising computational performance, can be made possible with the use of hashing techniques.

Hashing, converts the high-dimensional feature vectors of images into compact, low-dimensional hash codes. The primary goal of hashing lies in the creation of short memory efficient representations of the original images, ensuring that images that contain similar objects or patterns, will have more alike representations compared to those of dissimilar images. The original problem is then modified to the similarity of the hash codes that correspond to the dataset and query images, enabling efficient similarity assessment and retrieval.

To this end, our approach comprises two modules: the DL Indexer and the DL Comparator. The DL Indexer operates across multiple instances in a distributed manner in order to extract hashcodes from facial data provided by LEAs. The

DL Comparator is responsible for the suspect identification process. It systematically searches through all generated hash-codes to identify matches, generating a list of potential suspect matches from different LEAs.

The remaining paper is structured as follows: In Section II, prior works related to facial recognition frameworks that are developed to assist LEAs are discussed. Section III, provides a detailed description of our approach, which encompasses the functionality of both the DL Indexer and the DL Comparator, as well as the overall architecture. Section IV, presents the results of our approach, using an illustrative example to demonstrate its practical application. Finally, we conclude this paper in Section IV.

II. RELATED WORK

The utilization of emerging technologies in order to assist Law Enforcement Agencies (LEAs) in various operational tasks, has become increasingly prevalent. For many years, LEAs have relied on biometric data such as fingerprints and DNA for identifying and searching for potential matches. However, with the emergence of additional biometric data such as facial or gait information, new frameworks have been proposed to maximize the utility of these modalities for investigative and intelligence purposes.

Facial data's reduced discrimination power, poses challenges in operational deployment, database scalability, and distinguishing between investigative and evaluative use. For this, [1] has created a framework that emphasizes the need for interfaces facilitating visual comparisons and intelligent processing of forensic case data, thereby enhancing crime detection and monitoring.

In this direction, facial composites are commonly used by LEAs for suspect identification. However, existing facial composite to mugshot matching systems lack operational deployability. The FaceSketchID System [2] addresses this gap by offering a scalable and operationally deployable software solution that achieves state-of-the-art matching accuracy. By employing two diverse algorithms, the system matches facial composites to mugshot datasets, thus enhancing law enforcement capabilities.

The prevalence of social media websites has prompted LEAs to search for persons of interest among the billions of shared photos. Addressing this challenge, the proposed face search system offers a balance between accuracy and scalability on galleries with millions of images [3]. This system tackles the large-scale face search problem in the context of social media and other web applications, where face images are unconstrained in terms of pose, expression, and illumination. Utilizing Product Quantization (PQ), the system achieves efficient retrieval by first filtering faces and then re-ranking them using a commercial-off-the-shelf (COTS) face matcher. While similar to our approach in utilizing PQ, it employs a different model and datasets for experimentation.

Finally, in recent years there has been an approach to utilize smart glasses equipped with facial recognition capabilities [4], that offer advantages over traditional security cameras due

to their portability and ability to capture good frontal views. Unlike security cameras installed in streets, smart glasses can identify individuals involved in law-breaking activities and provide detailed information. This solution leverages facial recognition to achieve high-frequency and accurate recognition rates. By capturing images, matching them with a database, and displaying results on an OLED screen, these glasses enhance law enforcement capabilities.

While the reviewed literature showcases significant advancements in facial recognition technology, a notable gap persists in the interoperability of facial images across various systems and databases utilized by different LEAs. Despite the development of robust face search systems, facial composite matching software, and the integration of biometric modalities in forensic science, the lack of standardized protocols for sharing and processing facial data hinders effective collaboration and information exchange between agencies. Addressing this gap is crucial for enhancing the efficiency and accuracy of facial recognition systems in the pursuit of global security and law enforcement objectives.

III. THE APPROACH

In an ideal scenario where data privacy and security concerns are minimized, the collaboration between LEAs would require for the initiator LEA to send the facial image of a suspect to all collaborating LEAs requesting information. This would entail developing an efficient and robust methodology, focused on computational performance and accuracy.

Given the legal restrictions surrounding the provision of suspect information and the necessity for adherence to established legal procedures prior to any exchange of sensitive data, it is imperative to safeguard against any potential insight being gleaned by external agents from both the encoded data and its transmission.

To achieve this, a refined approach is employed, incorporating Product Quantization (PQ) for hashing and homomorphic encryption to ensure no data breaches occur. PQ involves the transformation of high-dimensional feature vectors representing image data, into concise and memory-efficient representations known as Product Quantization codes (PQ codes). This conversion simplifies the comparison of images into the comparison of PQ codes, which is more computationally efficient. Meanwhile, homomorphic encryption allows the conversion of data into ciphertext that can be analysed as if it were still in its original form. This encryption technique eliminates the need for decrypting data in order to compare their similarity, ensuring secure processing and sharing of information without compromising privacy.

The architecture of this approach can be found in Figure 1. Initially, the DL Indexer undergoes rigorous training using an annotated high-quality and relevant dataset. This training ensures that the module acquires the necessary intelligence to comprehend and categorize complex patterns within data. Once trained, the DL Indexer is deployed across multiple instances in a distributed manner. Its primary objective is to crawl and extract information from local authority data

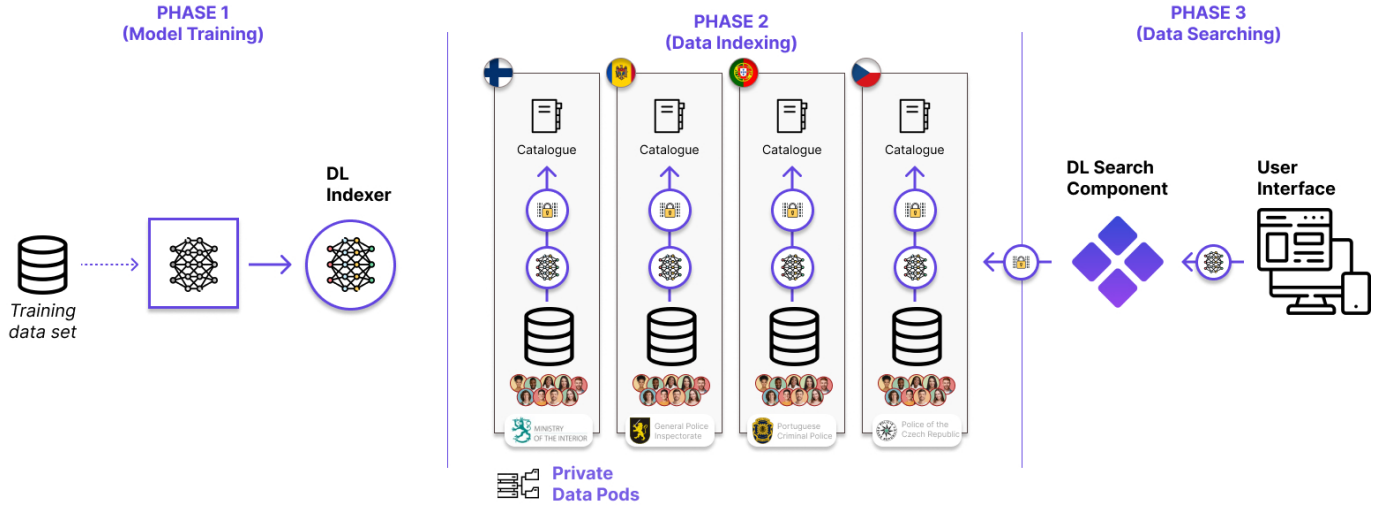


Fig. 1. Architecture

sources. This information is then transformed into a unique triplet format and then gets stored within a secure Data Space, ensuring secure and organized storage of crucial suspect-related information. Subsequently, the DL Comparator, upon receiving a query, accesses the global index. It employs advanced search algorithms to identify potential matches and returns a ranked list of triplets. This list provides insights into suspects that may span across different local authorities, enhancing the efficiency and accuracy of cross-authority suspect identification efforts.

A. Weight Sharing for Distributed Deployment

As mentioned above, the functionality of both the DL Indexer and DL Comparator relies heavily upon deep hashing methodologies. The neural network utilised by both methods is essential for creating the PQ codes and comparing two facial images. For this, we rely upon the Orthonormal product quantization network (OPQN) model introduced in [5]. Given the distributed deployment, each LEA possesses its own snapshot of the neural network. However, it is imperative to maintain consistency in inference results across these different deployments.

To achieve this consistency, global model weights are created and shared among all LEAs. These weights are trained using open-source datasets, namely FaceScrub and YouTube Faces, mitigating the risk of exposing proprietary information. Since these datasets are publicly available and devoid of private or restricted information, the original data remains secure. Moreover, this approach offers adaptability to dynamic environments, enabling seamless integration of new LEAs or updates to existing ones without compromising the integrity of the distributed framework.

The FaceScrub dataset [6] comprises a total of 106,863 face images of 530 celebrities, encompassing both male and female individuals. Each celebrity has about 200 images. In contrast, YouTube Faces dataset [7], contains 3,425 videos featuring 1,595 different people, with an average of 2.15 videos are

available for each subject. The average length of frames per video is approximately 181.3 frames, which are then utilized as the image data.

In accordance to Product Quantization, the global model weights create n disjoint subspaces, each of them consisting of 2^k centroids, also known as quantized vectors. The selection of the number of centroids is optimized for efficient memory usage by choosing a power of 2. As depicted in Figure 2, each centroid can be depicted as a vector with a designated Cluster ID.

This way each DL Indexer deployment will have a common ground when creating the PQ codes from the known suspects images and the comparison will be more reliable since the database of known suspects will be the same when created distributively as if it had been created by bringing all facial images in a central hub and then creating the PQ codes.

This approach also ensures that each DL Indexer deployment will have a common ground when creating the PQ codes from the known suspects' images, and the DL Comparator results will be more reliable since the database of known suspects will be the same when created distributively as if it had been created by bringing all facial images to a central hub and then creating the PQ codes. This method also eliminates the need for encryption since no data are exchanged.




Fig. 2. Global Model Weights

B. The DL Indexer

The functionality of the DL Indexer is illustrated in Figure 3. Initially, facial images acquired from LEAs undergo processing to generate feature vectors. These original feature vectors are then decomposed into disjoint subvectors, with each subvector corresponding to a specific subspace. In the next step, each centroid in the subspaces is compared against the corresponding feature sub-vector to determine the closest match. Each sub-vector is then encoded with the cluster ID of the closest centroid, which is often referred to as the reproduction value of the centroid.




Fig. 3. Feature Vector Quantization

As a result of splitting the original vector into n parts, the PQ code is encoded with n cluster ID numbers. For 2^k centroids, the memory required to store a PQ code is $n \cdot \log(k)$ bits. It's worth noting that Product Quantization is capable of producing a vast number of distinct distance values, estimated at $\binom{k}{2}^n$.

Finally, after the creation of the PQ codes, the information is transformed into triples with the following format: $\langle \text{hash_vector}, \text{LEA_ID}, \text{suspect_ID} \rangle$, where hash_vector is the produced PQ code, LEA_ID represents the local authority identifier and suspect_ID is an anonymised identifier that lacks individual meaning. However, each LEA retains a secure mapping that enables proper recognition.

C. The DL Comparator

The functionality of the DL Comparator begins with the provision of a suspect image. Batch processing is facilitated, ensuring consistent results whether provided with one image or multiple images. The suspect image, often a close-up picture of an arrested suspect, is referred to as the query image. This query image undergoes processing to generate the query feature vector.

This feature vector must be divided into n sub-vectors, following the same procedure used for the feature vectors generated by the DL Indexer. As depicted in Figure 4, each centroid in the subspaces is compared against the corresponding query feature sub-vector, to determine the distance between

them. This comparison results in an array where d_{ij} represents the distance between the i -th query sub-vector and the j -th centroid in subspace i .

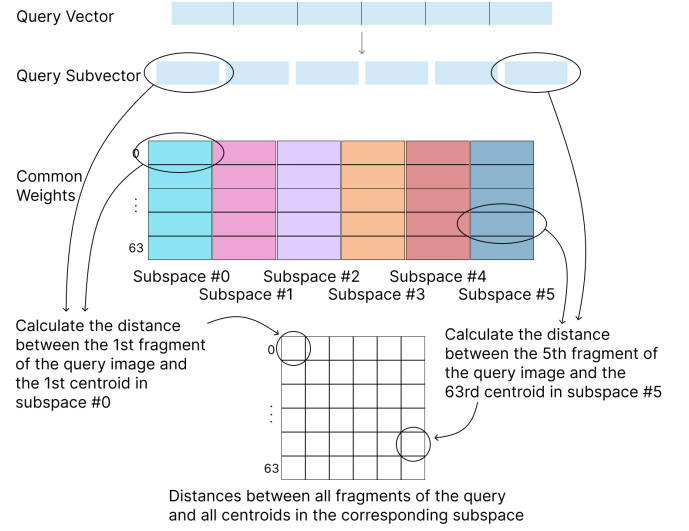


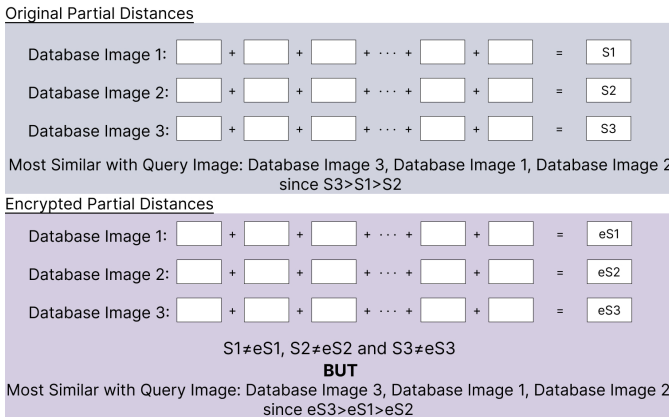
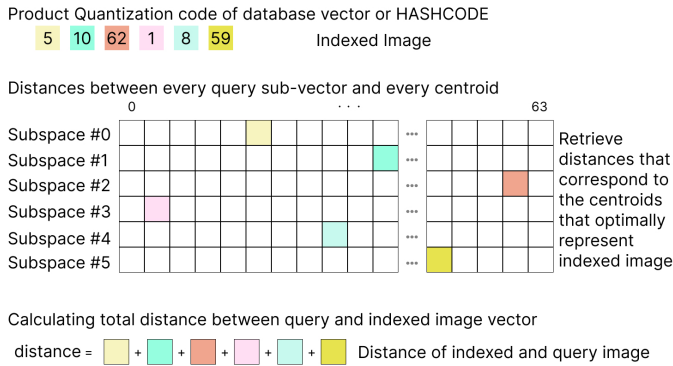
Fig. 4. Optimized Query Vector Centroid Matching

The next step involves collaboration between the LEAs, wherein the array of distances is shared with other agencies to identify potential matches. Despite the exchanged data being only numerical values, homomorphic encryption is employed to encrypt each of these values before sending them to the collaborating LEAs.

As shown in Figure 5, the cluster IDs that comprise the PQ codes are used as indexes for the array of distances. From the PQ codes we determine the centroids that are closest to the database image for each subspace. Instead of computing the actual distance between the query image and the database image, it is more efficient to calculate the difference between the query image and the closest centroids. While this distance is not exact, it provides a relatively close approximation at a lower computational cost.

Since the distance of the query image to each centroid is pre-calculated, the only computational requirement is to use the PQ codes as indexes to obtain the partial distances. The partial distance for subspace i is denoted as $d[i][PQcode[i]]$. The cumulative sum of these partial distances yields the final similarity distance between the query and database images.

It is worth noting the use of encryption and the fact that the above procedure is applied to encrypted data, without altering the overall result. Utilizing homomorphic encryption on the numerical values alters their individual values, and the encrypted and unencrypted results may not be identical. However, the saving grace of homomorphic encryption lies in the framework's reliance on the order of the results rather than their specific values. The crucial information is that the image of suspect A is more similar to the query image than the image of suspect B, rather than focusing on specific similarity degrees for each suspect. This aspect is elucidated in Figure 6



The final step of the DL Comparator is to combine and transform the results of all LEAs into triples with the following format: `<similarity_rate, LEA_ID, suspect_ID>`, where *similarity_rate* is one of the following options:

- Highly Likely Suspect
- Probable Suspect
- Unlikely Suspect
- Doubtful Suspect
- Highly Doubtful Suspect

while *LEA_ID* represents the local authority identifier and *suspect_ID* is an pseudonymised identifier that lacks individual meaning.

It is important to note that the returned IDs are pseudonymised, ensuring that the requesting LEA cannot obtain any knowledge regarding the true identities or facial data of the matches. Should further information be deemed necessary, LEAs will utilise alternative communication channels to exchange additional data, thereby preserving the confidentiality and integrity of the investigative process.

IV. RESULTS

To demonstrate the effectiveness of our approach, we conducted training of the model utilized by the DL Indexer and DL Comparator. The Global Weights shared by all snapshots

of the DL Indexer were generated by using a combination of the YouTube Faces datasets and 50% of the individuals from the FaceScrub Datasets, achieving a MAP of 89.34%.

The remaining 50% of the FaceScrub dataset was used to simulate the suspects apprehended by different LEAs. For each celebrity, 10% of their images were designated as query images to represent apprehended suspects for which each LEA lacks prior information and will seek assistance from others. The remaining 90% of images were allocated among the LEAs, reflecting the distribution of data in real-life scenarios. Notably, certain LEAs were intentionally provided with no information for specific celebrities, mirroring the diversity of information available to different agencies.

Each LEA uploads images of their apprehended suspects to their DL Indexer, allowing for the creation of PQ code representations. The DL Comparator utilizes this information to generate a list of potential matches, as showcased in Figure Figure 7. In the context of our solution there is no need for actual images of suspects to be depicted in a User Interface and it is utilized in this example solely for demonstrative purposes. The resulting IDs along with a similarity rate flag are provided and it is worth noting that these IDs hold no value for the initiator LEA, thus no additional information can be gained. Leveraging the similarity scale the initiator LEA can make informed decisions regarding potential collaboration with other LEAs in order to acquire vital information for the investigative process.

image data similarity and securely exchanging sensitive information. The distributed deployment of our approach ensures scalability and adaptability to dynamic environments, while adherence to data privacy regulations is maintained through pseudonymization and secure communication channels. We believe that our proposed solution will significantly contribute to improving the efficiency and effectiveness of international law enforcement efforts, ultimately enhancing global security and combating criminal activities. In the future, we aim to expand our approach beyond facial image data, to accept additional biometric data such as fingerprints and voice samples.

ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Union’s Horizon Europe research and innovation programme under the Grant Agreement No 101073920 (TENSOR). This publication reflects only the authors views. The European Union is not liable for any use that may be made of the information contained therein.

REFERENCES

- [1] D. Dessimoz and C. Champod, “A dedicated framework for weak biometrics in forensic science for investigation and intelligence purposes: The case of facial information,” *Security Journal*, vol. 29, pp. 603–617, 2016.
- [2] S. J. Klum, H. Han, B. F. Klare, and A. K. Jain, “The facesketchid system: Matching facial composites to mugshots,” *IEEE Transactions on Information Forensics and Security*, vol. 9, no. 12, pp. 2248–2263, 2014.
- [3] D. Wang, C. Otto, and A. K. Jain, “Face search at scale,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 6, pp. 1122–1136, 2016.
- [4] S. Khan, M. H. Javed, E. Ahmed, S. A. Shah, and S. U. Ali, “Facial recognition using convolutional neural networks and implementation on smart glasses,” in *2019 international conference on information science and communication technology (ICISCT)*. IEEE, 2019, pp. 1–6.
- [5] M. Zhang, X. Zhe, and H. Yan, “Orthonormal product quantization network for scalable face image retrieval,” *Pattern Recognition*, vol. 141, p. 109671, 2023.
- [6] H.-W. Ng and S. Winkler, “A data-driven approach to cleaning large face datasets,” in *2014 IEEE international conference on image processing (ICIP)*. IEEE, 2014, pp. 343–347.
- [7] L. Wolf, T. Hassner, and I. Maoz, “Face recognition in unconstrained videos with matched background similarity,” in *CVPR 2011*. IEEE, 2011, pp. 529–534.