

# Biometric Template Storage with Blockchain: A First Look into Cost and Performance Tradeoffs

Oscar Delgado-Mohatar, Julian Fierrez, Ruben Tolosana and Ruben Vera-Rodriguez  
Escuela Politecnica Superior  
Universidad Autonoma de Madrid, Madrid, Spain

{oscar.delgado, julian.fierrez, ruben.tolosana, ruben.vera}@uam.es

## Abstract

*We explore practical tradeoffs in blockchain-based biometric template storage. We first discuss opportunities and challenges in the integration of blockchain and biometrics, with emphasis in biometric template storage and protection, a key problem in biometrics still largely unsolved. Blockchain technologies provide excellent architectures and practical tools for securing and managing the sensitive and private data stored in biometric templates, but at a cost. We explore experimentally the key tradeoffs involved in that integration, namely: latency, processing time, economic cost, and biometric performance. We experimentally study those factors by implementing a smart contract on Ethereum for biometric template storage,<sup>1</sup> whose cost-performance is evaluated by varying the complexity of state-of-the-art schemes for face and handwritten signature biometrics. We report our experiments using popular benchmarks in biometrics research, including deep learning approaches and databases captured in the wild. As a result, we experimentally show that straightforward schemes for data storage in blockchain (i.e., direct and hash-based) may be prohibitive for biometric template storage using state-of-the-art biometric methods. A good cost-performance tradeoff is shown by using a blockchain approach based on Merkle trees.*

## 1. Introduction

The integration of the advantages and characteristics of public blockchains in biometric systems is a very recent area of research, but with a high potential and interest.

Combining blockchain and biometrics could potentially have many advantages. As a first approximation, the blockchain technology [20] could provide biometric systems with some desirable characteristics such as **immutability**, **accountability**, **availability** or **universal ac-**

<sup>1</sup>Deployed to the Ethereum Ropsten testnet at address: 0x8f737f448de451db9b1c046be7df3b48839673a1

**cess.** These properties enabled by blockchain technology may be very useful, among other applications in biometrics, to secure the biometric templates [14], and to assure privacy in biometric systems [3].

However, despite these opportunities, the current blockchain technology suffers from some potential limitations that must be carefully studied and characterized before the combination of both biometrics and blockchain technologies.

The main contribution of this study is two fold: 1) we analyze cost and performance tradeoffs when using blockchain for biometric template storage. We first discuss the existing alternatives for the storage of large volumes of data in blockchains, and how the complexity of schemes for face and handwritten signature biometrics affects to the cost and execution time of the final system; and 2) we experimentally measure these factors, optimizing the storage requirements of each biometric scheme while keeping their performances.

The remainder of the paper is organized as follows. In Section 2 a description of the most relevant features of blockchains, and the challenges and limitations of the technology that directly affect to biometric technologies is provided. In Section 3, we describe three popular storage techniques for public blockchains, briefly analyzing their main characteristics. Sections 4 and 5 present the setup and methods used in the experiments, whose results are shown in Section 6. Finally, Section 7 draws the final conclusions.

## 2. Blockchain for Biometrics

### 2.1. Smart contracts

A *smart contract* is, essentially, a piece of code executed in a secure environment that controls digital assets. Examples of these secure environments include regular servers controlled by “trusted parties”, decentralized networks (blockchains), or servers with secure hardware (SGX) [8, 9].

Many public blockchains support the execution of smart

contracts, but Ethereum [4] is currently considered the most reliable, secure and used. In essence, Ethereum could be seen as a distributed computer, with capability to execute programs written in Turing-complete, high-level programming languages. These programs comprise a collection of pre-defined instructions and data that has been recorded at a specific address of a blockchain. For biometric purposes, a smart contract running in a blockchain can assure a semantically correct execution.

## 2.2. Challenges and limitations

Despite the new opportunities already described in previous sections, the combination of both blockchain and biometric technologies is not straightforward due to the limitations of the current blockchain technology. Among them, it is important to remark: 1) its transaction processing capacity is currently very low (around tens of transactions per second), 2) its actual design implies that all system transactions must be stored, which makes the storage space necessary for its management to grow very quickly, and 3) its robustness against different types of attacks has not been sufficiently studied yet.

In addition, public blockchains suffer from other limitations which could impact the deployment and integration with biometric systems:

- **Economic cost of executing smart contracts:** In order to support smart contracts in blockchains (like Ethereum), and to reward the nodes that use their computing capacity to maintain the whole system, each instruction executed requires the payment of a fee in *gas* units. This gas is paid in the native cryptocurrency of Ethereum, called *ether*.
- **Privacy:** By design, all operations carried out in a public blockchain are known to all the participating nodes. Thus, it is not possible to directly use secret cryptographic keys, which reduces the number of potential applications.
- **Processing capability:** Another important limitation is related to its processing capability. Ethereum, for example, is able to run just around a dozen transactions per second, what it could be not enough for some scenarios.
- **Scalability:** Currently, the size of the public blockchains (Bitcoin and Ethereum) is around 200GB, and it is growing very fast. This can be a problem for some application scenarios such as the Internet of Things (IoT).

## 3. Storage requirements analysis

As stated in the previous section, one of the main potential limitations for the integration of both technologies is the

Operation	Gas/KB	ETH/KB	\$/KB
READ	6,400	0.000032	\$0.004
WRITE	640,000	0.0032	\$0.448

Table 1. Non-volatile storage costs in Ethereum. We have considered a gas price of 1 gwei (1 gwei =  $10^{-9}$  ETH), and 1 ETH = \$140 (at time of writing, March 2019).

cost of running a biometric system (totally or partially) in a blockchain. It is therefore crucial to properly estimate and minimize that cost. The present paper is an initial attempt in that regard.

This section describes the different existing schemes to store large volumes of data (e.g., a database of biometric templates) in public blockchains with smart contracts execution capabilities, like Ethereum.

There are essentially three approaches, which are presented below in terms of complexity (from lower to higher), and economic cost (from higher to lower):

- **Full on-chain storage:** all data is stored, as-is, in the blockchain.
- **Data hashing:** the blockchain only stores a hash of the data that guarantees its immutability. The data itself is stored off-chain in other system: distributed (e.g., IPFS [1]), cloud, or local.
- **Merkle trees:** data is stored also off-chain, but it is preprocessed by constructing a Merkle tree structure, which reduces storage costs and increases the bandwidth.

These alternatives are discussed in more detail next.

### 3.1. Full on-chain storage

This is the simplest scheme and therefore, the most inefficient and costly. In this case, the data are just stored in the blockchain as is, without any type of pre-processing. For example, biometric templates could be directly stored as a data structure in a smart contract, as part of a more general digital identity model.

In general terms, the storage space in public blockchains is specially expensive compared to computation, in order to discourage its abusive use. Therefore, as shown by experiments and figures presented in Section 6, the use of this storage scheme would commonly imply a prohibitive cost for most biometric applications.

As an example, Table 1 depicts the cost of reading and storing 1 Kilobyte of data in Ethereum in terms of gas units, ether, and US dollars.

### 3.2. Data hashing

To overcome the problems of the previous scheme, a more efficient approach is to store the data *off-chain* and

use the blockchain just as a integrity guarantee due to its intrinsic immutability. This way, instead of the full data, only a hash value of it is stored in the blockchain (smart contract). Then, the complete template can be stored in any other traditional external storage system (see Figure 1).

This possibility provides a great flexibility, because any platform, as public clouds or existing corporate servers, can be used to store the full set of biometric templates. In any case, to maintain the distributed spirit, resistance to censorship and high availability of public blockchains, distributed storage systems such as IPFS or Swarm [16] would be desirable in this case.

On the other hand, this approach can make use of any cryptographic hash function, such as the SHA3 family, which can produce outputs from 224 to 512 bits in length [2]. In this work, we consider hashes of 256 bits per template, which, in any case, can greatly reduce storage costs compared to full on-chain storage.

One drawback of this approach is that it is still necessary to ensure the availability of the data stored outside the blockchain. If these data were lost or tampered, even when this modification would be always noticed, the viability of the system would be compromised.

### 3.3. Merkle trees

Finally, the previous scheme can be still further improved, through the use of a data structure known as *Merkle tree* [12]. This construction is widely used in cryptography and computer science problems such as database integrity verification [13], peer-to-peer networks [18] and, of course, blockchains [4].

A Merkle tree is a binary tree data structure in which every node contains the cryptographic hash of the concatenation of its child nodes contents. Due to this recursive way of constructing itself, the tree root contains statistical information of the rest of nodes, and the modification of any node content will cause the complete change of the value of the root. This way, the integrity of an arbitrary amount of data can be efficiently assured by arranging this data in a Merkle tree form and securely storing the contents of its root node.

Regarding biometric template protection using blockchains, a biometric system using this technique would maintain a Merkle tree, storing a template at each node and assuring the root node in a smart contract. Therefore, when a new biometric template is created (after the enrollement stage), or an existing one is modified or deleted, the tree is re-calculated and the new root is updated in the blockchain. A simplified scheme of this approach can be found in Fig. 1 (right).

## 4. Experimental methods

### 4.1. Blockchain technology

The baseline architecture considered for performing the experiments presented in this work was initially introduced in [5]. This architecture substitutes the usual template database of a biometric system by a blockchain, adding basic operations (i.e., creation, modification and deletion of templates) through the use of smart contracts.

This design provides some advantages:

- The modifications to the existing biometric architectures are minimal, so that usual biometric techniques and algorithms (e.g., feature extraction and matching) can be used normally.
- Since the biometric process is performed off-chain, this architecture avoids the scalability problems of public blockchains (except in a massive batch of user registration during the system startup, for example).
- No need to use complex smart contracts, which facilitates development and reduces execution costs. Smart contracts do not implement biometric “logic”, but only the minimum necessary functions to manage the storage of the templates.

As stated, we have implemented a basic smart contract, that has been deployed to the Ropsten Ethereum testnet. The contract models a biometric template as a data structure `BiometricTemplate` implemented as a raw array of bytes. This structure is stored in a mapping (or hash table), with an identifier number for the user acting as the mapping key `mapping(uint => BiometricTemplate)`. The source code of smart contracts can be found in Appendix A. The main operations are described below:

- **Creation:** Receives the user ID, template data and metadata, and adds a new `BiometricTemplate` structure to the blockchain.
- **Modification:** Modifies the template of an existing user. For a hash table storage scheme, this is equivalent to an addition operation.
- **Deletion:** Removes the link between a specific template and user ID. Due to the public nature of Ethereum, technically the old template data remains forever in the blockchain.
- **Retrieval:** Retrieves the `BiometricTemplate` structure for a user. This function is a *call*, not a transaction as the rest of functions. This operation is usually read-only (and, therefore, free to execute), while the previous three operations were potentially state-changing.

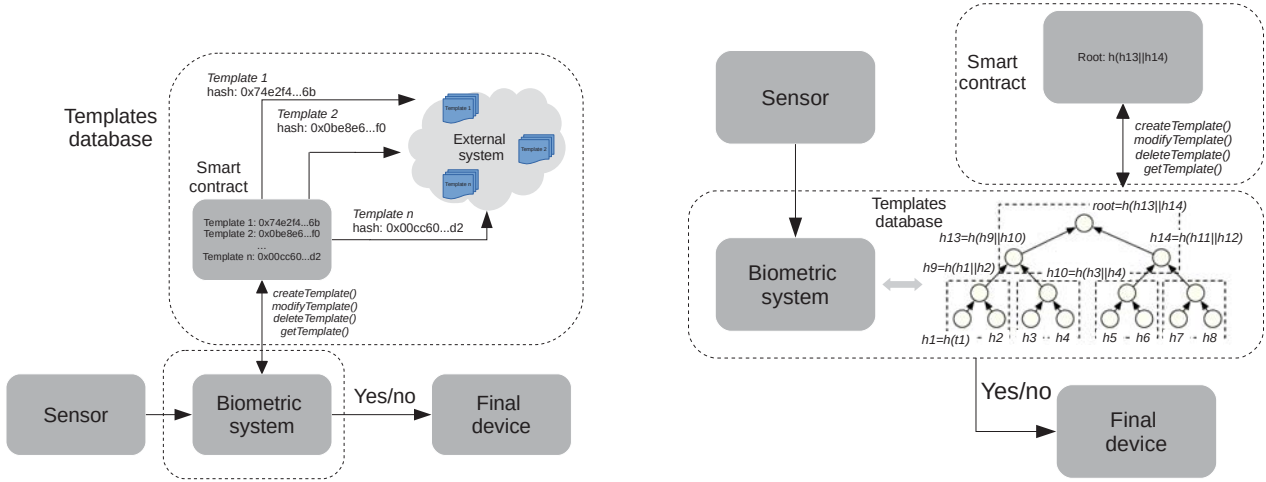


Figure 1. Biometric systems using data hashing (left) and Merkle trees (right) blockchain storage techniques.

## 4.2. Biometric systems

Two different biometric traits are considered in the analysis: 1) face, and 2) dynamic signature. In this way we experiment both with image-based physiological biometrics, and with signal-based behavioral biometrics.

### 4.2.1 Face biometrics

One of the most popular face recognition based on deep convolutional neural networks (DCNNs) are evaluated in this study: VGG-Face [17].

In this system images are propagated through the CNN obtaining the features at the last fully connected layer. The final matching score is computed through the Euclidean distance of the features obtained from each face image. The dimension of the face features are of 4,096.

### 4.2.2 Dynamic signature biometrics

Two popular approaches are evaluated in this study: *i*) feature-based systems (a.k.a. global systems), and *ii*) time functions-based systems (a.k.a. local systems).

For the global system, we extract for each signature a total of 100 global features from the normalized  $X$  and  $Y$  spatial coordinates. These features are described in [11], and are related to time, kinematic, direction, and geometry information. For the similarity computation, the Mahalanobis distance is used to compare the similarity between a signature and a claimed user model.

For the local system, a total of 21 local features are extracted from the normalised signals  $X$  and  $Y$  spatial coordinates [19]. For the similarity computation, DTW is used to compare the similarity between genuine and query input samples, finding the optimal elastic match among time sequences that minimises a given distance measure.

## 5. Experimental protocol

This section describes the main characteristics of the biometric systems evaluated, and the key tradeoffs involved in the integration of blockchain technology in both *face* and *dynamic signature* biometric traits.

This integration is evaluated in terms of cost of storage, execution, and performance in the Ethereum blockchain.

Two popular biometric databases are considered for the analysis of face and signature biometrics: Labeled Faces in the Wild (LFW) [10], and Biosecure [15].

### 5.1. Face

#### 5.1.1 Databases

Face verification experiments are conducted on the LFW database (Labeled Faced in the Wild) [7]. LFW is one of the most popular datasets used in face recognition with more than 13,000 face images of famous people collected from the web. We have used the aligned dataset where each image was aligned with funneling techniques.

#### 5.1.2 Data processing

The VGG-Face pre-trained model was tested using the *unrestricted* and *outside training data* protocols proposed in [7]. The VGG-Face model was trained with VGG-Face database [17], therefore there is not extra training for the pre-trained model used here. The evaluation results are computed for 6,000 one-to-one comparisons composed by 3,000 genuine pairs (pairs of images from the same person) and 3,000 impostor pairs (pairs of images belonging to different persons) following the protocols from LFW database.

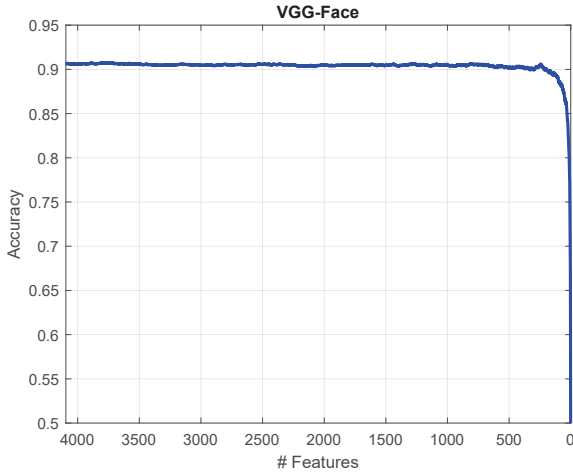


Figure 2. System performance results in terms of the size of the feature embeddings for VGG-Face CNNs model.

## 5.2. Dynamic signature

### 5.2.1 Databases

The dynamic signature verification technology is analyzed using the Biosecure DS2 dataset. This dataset was captured using a Wacom Intuous 3 digitizing tablet with an inking pen in an office-like scenario, providing the following information:  $X$  and  $Y$  spatial coordinates, pressure, and timestamp (sampling frequency 100 Hz).

### 5.2.2 Data processing

In this study, we consider a set of 50 users. For each user, the first 5 genuine signatures of the first session are used for training, whereas the 15 genuine signatures of the second session are left for testing in order to consider the inter-session variability. In this study we analyze the robustness of our proposed system against random (zero-effort) forgeries. Scores are obtained by comparing the training signatures with one genuine signature of the remaining users. For the global system, scores are obtained by comparing signatures against the user model, while for the local system, the average score of the five one-to-one comparisons is used.

## 5.3. Blockchain integration tradeoffs

### 5.3.1 Face

The system performance results in terms of EER (%) for VGG-Face CNN model is depicted in Fig. 2 for different sizes of the biometric template. This analysis has been carried out by removing features randomly from the original feature embedding.

Analyzing results in Fig. 2, in general the system performance is very stable while we gradually remove features.

VGG-Face is able to obtain a verification rate with an accuracy of 89% only using 100 features (only 2.5% from the original 4096 features). This behavior shows that there is a very high redundancy within the feature embedding of CNNs face models, which makes possible to obtain very competitive verification performance while keeping only a small set of features.

### 5.3.2 Dynamic signature

The system performance results in terms of EER (%) of both global and local systems are depicted in Fig. 3 for different sizes of the biometric template. This analysis has been carried out using Sequential Forward Floating Search (SFFS) in order to select the best subsets of global and local features that improve the system performance in terms of EER (%).

Analyzing in Fig. 3 (left) the global approach, the system performance improves when increasing from 1 to 30-40 global features. After that, a degradation of the system performance is produced when adding more global features to the optimal feature vector. Therefore, in order to reduce the cost of saving the biometric templates in the blockchain platform, and also achieve the best possible system performance, we propose to save the best 30 global features in the biometric template, achieving this way a final 1.5% EER.

The same analysis has been carried out for the local approach in Fig. 3 (right). The system performance improves when adding more local features, achieving for the best system performance a final 0.5% EER using 9 local functions (total number of features = 9 local functions  $\times$  average time samples per signature = 3,087 features).

## 6. Experimental results

This section analyzes the results of the evaluation of the integration of the biometric systems previously described in Ethereum, resumed in Table 2.

The smart contract developed has been written in Solidity language, and deployed to the Ethereum Ropsten testnet at the address `0x8f737f448de451db9b1c046be7df3b48839673a1`, where can be verified with any blockchain explorer like Etherscan [6]. It is a basic contract, which has not been optimised and does not take care of security issues, and should be used only for experimental purposes.

Table 2 shows the costs of the different operations over the templates (creation, modification, deletion, and retrieval) in units of gas and US dollars, for the biometric technologies and blockchain storage schemes evaluated.

The results clearly prove that the most efficient storage scheme is the one based on Merkle trees. In fact, it is the only one capable of storing any amount of data for the same






Figure 3. System performance results in terms of the size of the optimal feature/time function vector selected by the SFFS algorithm. Left: global system. Right: local system. For local system: #Features = #Local Functions  $\times$  Average time samples per signature (=343).

Biometric		Operation	Storage scheme			Performance
Scheme	Template size		Full on-chain	Data hashing (cost per template)	Merkle trees (cost for any number of templates)	Execution time (average)
-	-	Smart contract deployment	498274 gas (\$0.06972)			19.19 secs
Signature	Global 30 x 16 bits	Creation	108844 gas (\$0.014)	86848 gas (\$0.0122)		10.66 secs
		Modification	21378 gas (\$0.003)	18850 gas (\$0.0026)		11.55 secs
		Deletion	-	-	-	-
		Retrieval	-	-	-	-
	Local 3087 x 16 bits	Creation	4358990 gas (\$0.610)	86848 gas (\$0.0122)		12.61 secs
		Modification	504322 gas (\$0.07)	18850 gas (\$0.0026)		12.85 secs
		Deletion	-	-	-	-
		Retrieval	-	-	-	-
Face	VGG-Face 100 x 32 bits	Creation	352912 gas (\$0.049)	86848 gas (\$0.0122)		10.53 secs
		Modification	49192 gas (\$0.0068)	18850 gas (\$0.0026)		16.38 secs
		Deletion	-	-	-	-
		Retrieval	-	-	-	-

Table 2. We have considered a gas price of 1 gwei (1 gwei =  $10^{-9}$  ETH), and 1 ETH = \$140 (accurate at time of writing, March 2019).

cost. The rest of the schemes would quickly have a prohibitive cost for the number of templates to be stored in a real environment.

For example, protecting a million of templates would cost between \$14,000 and \$610,000 for the signature system, and \$49,000 for VGG Face using the *full on-chain* storage scheme. Clearly this is not a realistic option, discouraged not only in economic terms, but also for security and performance reasons.

The *data hashing* scheme would improve significantly those figures, because it does not store the data itself, but only a hash that guarantees the integrity. For the same sce-

nario, the cost would be a much more reasonable amount of \$12,200 for all the biometric technologies.

Finally, the *Merkle trees* scheme would imply a cost of only one cent of dollar (\$0.0122) for the storage of any amount of templates. In addition, also the modification operation of a template would have the same cost. However, even for a biometric system operating in a large corporation or environment, these costs seem reasonable.

Of course, all these prices could vary greatly depending on the price of ether, which, as the rest of cryptocurrencies, usually suffers sharp increases and falls in price. However, because it only needs to store 256 bits regardless of the total

volume of data, the Merkle tree scheme would still have a reasonable cost in any case.

In terms of execution time and performance, the experiments also show that this hybrid system is viable. It is important to note that the tests have been carried out in a testnet, where the confirmation times are higher and have greater variability than in the mainnet. Times have been measured performing each operation ten times, discarding the minimum and maximum times, and calculating the average of the rest.

As can be seen, the execution time is slightly higher than 10 seconds for most of the operations and biometric systems, which seems an acceptable time for the usability of the system even during a user enrollment, for example.

Finally, the retrieval operation, necessary for the verification of a template, is a read-only operation and, therefore, free of cost. In addition, it can be also considered immediate in terms of execution time, due to that the request is processed by the local Ethereum node, and it does not reach the network.

## 7. Conclusions

In this paper we have explored the viability of biometric systems based on blockchain with focus on storing the biometric templates. This experimental exploration has been around key cost-performance tradeoffs, in particular: time of execution of the transactions, economic cost, and biometric performance.

We have first discussed the main storage schemes for public blockchains (Ethereum), and implemented a smart contract for the estimation of its storage cost. The results obtained prove that straightforward schemes such as the direct storage of the biometric templates on-chain, or direct data hashing, are not appropriate for a real biometric system. However, when Merkle trees are included as an intermediate data structure, the storage costs become fixed regardless the total volume of data to store, and reduced execution times (between 10 - 20 seconds for *write* operations) are obtained. The *read* operations (retrieving) of templates are usually free of cost and very fast to execute, because they are processed locally.

In brief, in this work we have shown that the integration of biometric and public blockchains is possible both from an economic and performance perspective, including two case studies with state-of-the-art methods and protocols in face and signature biometrics.

## Acknowledgments

This work has been supported by projects: BIBECA (TEC2018 MINECO), Bio-Guard (Ayudas Fundacin BBVA a Equipos de Investigacin Cientfica 2017), UAM Cecabank chair on Biometrics, and UAM-Grant Thornton chair on

Blockchain. Ruben Tolosana is supported by a FPU Fellowship from Spanish MECD.

## References

- [1] Juan Benet. IPFS - Content Addressed, Versioned, P2P File System. jul 2014. [2](#)
- [2] Guido Bertoni, Joan Daemen, Michaël Peeters, and Gilles Van Assche. Keccak. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology (EUROCRYPT)*, pages 313–314. Springer, 2013. [3](#)
- [3] J. Bringer, H. Chabanne, and A. Patey. Privacy-preserving biometric identification using secure multiparty computation: An overview and recent trends. *IEEE Signal Processing Magazine*, 30(2):42–52, March 2013. [1](#)
- [4] Chris Dannen. *Introducing Ethereum and Solidity: Foundations of Cryptocurrency and Blockchain Programming for Beginners*. Apress, Berkeley, CA, USA, 2017. [2](#), [3](#)
- [5] Oscar Delgado-Mohatar, Julian Fierrez, Ruben Tolosana, and Ruben Vera-Rodriguez. Blockchain and biometrics: A first look into opportunities and challenges. *arXiv e-prints*, page arXiv:1903.05496, Mar 2019. [3](#)
- [6] Oscar Delgado-Mohatar et al. Smart contract source code. <https://ropsten.etherscan.io/address/0x8f737f448de451db9b1c046be7df3b48839673a1/>, 2019. [Online; accessed 22-March-2019]. [5](#)
- [7] Gary B. Huang, Marwan Mattar, Tamara Berg, and Eric Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. In *Proc. Workshop on Faces in 'Real-Life' Images: Detection, Alignment, and Recognition*, Marseille, France, Oct. 2008. [4](#)
- [8] Vishal Karande et al. SGX-Log: Securing system logs with SGX. In *Proc. of ACM Asian Conf. on Computer and Communications Security*, 2017. [1](#)
- [9] Kubilay A. Küçük et al. Exploring the use of Intel SGX for secure many-party applications. In *Workshop on System Software for Trusted Execution*, 2016. [1](#)
- [10] Erik Learned-Miller, Gary B Huang, Aruni RoyChowdhury, Haoxiang Li, and Gang Hua. Labeled faces in the wild: A survey. In *Advances in face detection and facial image analysis*, pages 189–248. Springer, 2016. [4](#)
- [11] Marcos Martinez-Diaz, Julian Fierrez, and Seiichiro Hangai. Signature features. In Stan Z. Li and Anil K. Jain, editors, *Encyclopedia of Biometrics*, pages 1375–1382. Springer, 2015. [4](#)
- [12] Ralph C. Merkle. A digital signature based on a conventional encryption function. In *Conf. on the Theory and Applications of Cryptographic Techniques on Advances in Cryptology (CRYPTO)*, pages 369–378, London, UK, 1988. Springer-Verlag. [3](#)
- [13] Kyriakos Mouratidis, Dimitris Sacharidis, and Hweehwa Pang. Partially materialized digest scheme: An efficient verification method for outsourced databases. *The VLDB Journal*, 18(1):363–381, Jan. 2009. [3](#)
- [14] K. Nandakumar and A. K. Jain. Biometric template protection: Bridging the performance gap between theory and

- practice. *IEEE Signal Processing Magazine*, 32(5):88–100, Sep. 2015. 1
- [15] Javier Ortega-Garcia, Julian Fierrez, et al. The Multi-scenario Multienvironment Biosecure Multimodal Database (BMDB). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1097–1111, 2010. 4
- [16] Kazim Rifat Ozyilmaz and Arda Yurdakul. Designing a blockchain-based iot infrastructure with ethereum, swarm and lora. *CoRR*, abs/1809.07655, 2018. 3
- [17] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proc. British Machine Vision Conference*, 2015. 4
- [18] H. B. Ribeiro and E. Anceaume. Datacube: A p2p persistent data storage architecture based on hybrid redundancy schema. In *2010 18th Euromicro Conference on Parallel, Distributed and Network-based Processing*, pages 302–306, Feb 2010. 3
- [19] Ruben Tolosana, Ruben Vera-Rodriguez, Javier Ortega-Garcia, and Julian Fierrez. Preprocessing and feature selection for improved sensor interoperability in online biometric signature verification. *IEEE Access*, 3:478–489, 2015. 4
- [20] Wenbin Zhang, Yuan Yuan, Yanyan Hu, Karthik Nandakumar, Anuj Chopra, Sam Sim, and Angelo De Caro. Blockchain-based distributed compliance in multinational corporations’ cross-border intercompany transactions. In Kohei Arai, Supriya Kapoor, and Rahul Bhatia, editors, *Advances in Information and Communication Networks*, pages 304–320. Springer, 2019. 1



## A. Appendix: Smart contract source code

```
pragma solidity >=0.4.22 <0.6.0;
contract BioBlockchain {

    /// This struct models a simple biometric template
    struct BiometricTemplate {
        bytes templateMetadata;
        bytes templateData;
    }

    /// Each template is indexed by an user ID
    mapping(uint => BiometricTemplate) templates;

    /// Store a new template
    function createNewTemplate(uint _templateID,
                               bytes memory _templateMetaData,
                               bytes memory _templateData) public {

        /// Add new template to mapping
        templates[_templateID].templateMetadata = _templateMetaData;
        templates[_templateID].templateData = _templateData;

    }

    /// Return a user template
    function getTemplate(uint _templateID) view public returns (bytes memory) {

        return(templates[_templateID].templateData);

    }

    /// Modify a user template
    function modifyTemplate(uint _userID,
                            bytes memory _newTemplateMetaData,
                            bytes memory _newTemplateData) public {

        // Due to that a mapping is internally implemented using
        // a hash table, the modification operation is equivalent
        // to a insertion
        createNewTemplate(_userID, _newTemplateMetaData, _newTemplateData);

    }

    /// Return an specific template
    function deleteTemplate(uint _userID) public {

        delete templates[_userID];

    }

}
```