

Inference-Based Similarity Search in Randomized Montgomery Domains for Privacy-Preserving Biometric Identification

Yi Wang[✉], Jianwu Wan, Jun Guo, Yiu-Ming Cheung, and Pong C. Yuen[✉]

Abstract—Similarity search is essential to many important applications and often involves searching at scale on high-dimensional data based on their similarity to a query. In biometric applications, recent vulnerability studies have shown that adversarial machine learning can compromise biometric recognition systems by exploiting the biometric similarity information. Existing methods for biometric privacy protection are in general based on pairwise matching of secured biometric templates and have inherent limitations in search efficiency and scalability. In this paper, we propose an inference-based framework for privacy-preserving similarity search in Hamming space. Our approach builds on an obfuscated distance measure that can conceal Hamming distance in a dynamic interval. Such a mechanism enables us to systematically design statistically reliable methods for retrieving most likely candidates without knowing the exact distance values. We further propose to apply Montgomery multiplication for generating search indexes that can withstand adversarial similarity analysis, and show that information leakage in randomized Montgomery domains can be made negligibly small. Our experiments on public biometric datasets demonstrate that the inference-based approach can achieve a search accuracy close to the best performance possible with secure computation methods, but the associated cost is reduced by orders of magnitude compared to cryptographic primitives.

Index Terms—Biometric identification, privacy protection, nearest neighbour search, hypothesis testing, multi-index hashing

1 INTRODUCTION

BEYOND conventional applications in forensic investigations and border control, personal identification based on biometrics (e.g., face, iris, fingerprint) is also expanding in the private sector. For instance, hundreds of US health-care organizations are using iris and palm-vein biometrics for patient identification and authentication [1]. India's DCB Bank has introduced Aadhaar-based fingerprint-reading ATMs where customers can access their bank accounts using their Aadhaar number and biometric details instead of a PIN [2]. The service is required to connect with the Aadhaar server for authenticating the identity of a customer every time a transaction is initiated. Today's social media websites provide photo-sharing services and are expected to handle face search in a database on the order of millions of enrolled records every day [3]. Such applications of biometric identification in general involve handling queries and searches at scale in a networked environment.

Biometric identification systems typically require *one-to-many* comparisons by evaluating *biometric similarity* between an input query and the database records in some retrieval space. Security analysis has shown that information leakage in storing and processing biometric data can lead to identity theft and adversarial tracking [4], [5]. In particular, *adversarial machine learning* of biometric recognition is found possible in the retrieval space [6]. Even if individual templates are secured in the database, it was shown that an adversary can still exploit the biometric similarity information to compromise system operations. For instance, the similarity information can be used to fabricate a biometric spoof in a hill-climbing attack [7] or to perform a biometric equivalent of the dictionary attack [8]. Another example is in multibiometric systems where multiple biometric traits are combined to increase the recognition accuracy and the population coverage of large-scale personal identification. Without fabricating any fake trait, it is possible to evade multibiometric systems if the attacker seizes the genuine similarity distribution of one single trait [9], [10]. Security and privacy risks associated with the biometric similarity information signify the importance of protecting it in the process of biometric identification.

Binary representations of biometric features are of growing interest for search applications as they can ensure fast matching operations in Hamming space [11], [12], [13]. To achieve high discriminative power and matching accuracy, binary representations of biometric features typically contain hundreds, if not thousands, of bits [13], [14]. Moreover, due to acquisition noise and other factors, biometric features are known to have inherently large intra-class variations between samples that are taken from the same trait.

- Y. Wang and J. Guo are with the School of Computer Science and Network Security, Dongguan University of Technology, Dongguan 523808, China. E-mail: {wangyi, guojun}@dgut.edu.cn.
- J. Wan is with the Department of Computer Science, Changzhou University, Changzhou 213164, China. E-mail: jianwuwang@gmail.com.
- Y.-M. Cheung and P.C. Yuen are with the Department of Computer Science, Hong Kong Baptist University, Hong Kong SAR, China. E-mail: {ymc, pcyuen}@comp.hkbu.edu.hk.

Manuscript received 13 Oct. 2016; revised 25 May 2017; accepted 1 July 2017.
Date of publication 13 July 2017; date of current version 12 June 2018.
(Corresponding author: Jun Guo.)

Recommended for acceptance by H. Ling.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2017.2727048

This indicates that the input query is not likely to find an identical match in the biometric database and the true match may have a relatively large *distance*, i.e., the number of mismatching bits, to the query for similarity evaluation in Hamming space. Rather than exhaustively exploring the database by pairwise comparisons, nearest neighbour (NN) methods for long binary representations with a large search radius of interest are proposed for fast retrieval in Hamming space [11], [15], [16]. These methods are designed to reduce the matching complexity of a similarity search by using certain data structures, e.g., hash tables.

Without privacy protection in design, NN methods have built-in mechanisms that are vulnerable to information leakage. For example, search indexes are often generated directly from biometric features through *distance-preserving* transformations [12], [13], [17], [18]. This allows an adversary who can access the data structures to analyse the search indexes and glean critical similarity information of the dataset. Another source of information leakage is at the retrieval stage. To increase the retrieval accuracy, NN methods generally require distance computations and comparisons between the query and matching candidates retrieved from the data structure. From the perspective of privacy protection, it is desirable to return the retrieval results in order without revealing their distance values [19].

Secure multiparty computation techniques are proposed in the context of *one-to-one* verification to protect biometric similarity evaluations [5], [19], [20], [21]. They are mostly based on combining cryptographic primitives such as homomorphic encryption and garbled circuits to process encrypted biometric data without the need of decrypting them. As distance measures are not preserved in the encrypted domain, adversarial learning can be effectively prevented. In our context of one-to-many identification, however, a similarity search must return multiple candidates in order according to their similarity measures to the query. The process involves not only intensive distance computations but also intensive distance comparisons that are especially cumbersome to perform in the encrypted domain [22].

In this paper, we propose an inference-based framework for privacy-preserving similarity search of biometric identities in Hamming space. We argue that, from the perspective of adversarial machine learning, an attacker must collect sufficient and reliable training data that characterise the underlying biometric feature space. This motivates us to design effective mechanisms to diffuse the feature information over piecewise obfuscated sub-hash codes, rather than concealing everything in the encrypted domain. In this way, we aim to facilitate information-theoretic privacy in the design of hash-based indexing methods without sacrificing too much of search accuracy and efficiency.

The main contributions of this paper are:

- We start with the Hamming-ball search problem and model it as binary channel estimation for which we derive the optimal decision rule based on the theory of binary hypothesis testing. This enables the design of an inference-based approach for detecting, with a high accuracy in the retrieval space, matching candidates within a predefined neighbourhood of the query without knowing the actual distance values.

- As a key element of the binary channel, we present a distance obfuscation mechanism based on matching piecewise obfuscated binary codes for hash-based similarity search. We derive analytically the condition where it is guaranteed to have a substring collision for neighbouring pairs, and show how it gives rise to an obfuscated distance measure that conceals the Hamming distance in a dynamic interval.
- We show how the inference-based approach for Hamming-ball search can be extended to handle a test-based rank-ordered search that arranges the retrieved candidates in rank order without the need of comparing distance values. We further show how to perform an approximate NN search by choosing candidates based on the obfuscated distance measure with or without the need of test-based ranking.
- We apply Montgomery multiplication to generate privacy-preserving search indexes that allow detection of substring collisions in randomized Montgomery domains. This enables the design of a search protocol that restricts all parties involved in the process to perform their respective computations in separate Montgomery domains for privacy protection.
- We study the privacy-preserving strength of our search scheme in an information-theoretic approach. In particular, we quantify the privacy gain of indexing in randomized Montgomery domains, and analyze how difficult it is for an adversary to infer biometric identities from search indexes in the proposed indexing and retrieval process.

The remainder of this paper is organized as follows. Section 2 reviews the related work. Section 3 provides the details of the inference-based similarity search framework. Section 4 presents the design of index generation in randomized Montgomery domains and the description of the search protocol. Analysis of the privacy-preserving strength is presented in Section 5. Performance evaluation is reported in Section 6. Finally, we draw conclusions in Section 7.

2 RELATED WORK

Biometric privacy study has been largely focused on enabling one-to-one matching without revealing the biometric features that characterize an individual, known as biometric template protection [4], [23], [24]. Depending on how the protected reference is generated and matched, template protection schemes can be classified into bio-cryptosystems and feature transformations [4]. Bio-cryptosystems generate error correction codewords for an indirect matching of biometric templates [23]. They typically yield a yes/no decision for verification on one-to-one basis, which is not suitable for returning candidates in the order of their matching degrees. Feature transformation methods apply non-invertible functions to biometric templates [24]. The protected references are usually distance preserving to perform matching directly in the transformed space.

Recently, secure signal processing methods have been introduced to protect biometric matching by concealing both the feature contents and similarity evaluation in the encrypted domain [5], [20], [25], [26]. These techniques use cryptographic primitives as a wrapper of distance-related

functions. Thus, a privacy-preserving similarity search can be decomposed into two steps: secure distance computation followed by minimum (distance) finding via oblivious transfer [27]. As unauthorized information gathering is precluded in the encrypted domain, secrecy and accuracy can be ensured in pairwise matching for biometric identification. However, cryptographic primitives are generally associated with high computation costs and excessive communication overheads. For instance, minimum finding is very complex to perform in the encrypted domain. It demands pairwise comparison of all encrypted distance values and frequent interactions between the query and the database [22]. This makes it intrinsically difficult, if not impossible, to meet the efficiency and scalability requirements of search applications, especially when dealing with high-dimensional data, in the encrypted domain [22], [27].

Other than search in the encrypted domain, the concept of search with reduced reference is proposed in privacy-preserving content-based information retrieval to protect the original content and accelerate the search simultaneously [22], [28], [29]. The basic idea therein is to enforce k -anonymity or l -diversity properties by raising the ambiguity level of a data collection [30], [31]. Techniques of this paradigm are mostly based on randomized embedding [27] and in particular locality sensitive hashing (LSH) [28], [32], [33], [34]. LSH performs approximate NN search by hashing similar items into the same bucket and those that are distant from one another into different buckets, respectively, with high probabilities. However, LSH by itself does not guarantee privacy [27], [33], [34]. It requires all parties involved in a search to use the same random keys in generating hash codes. Moreover, to achieve a good precision in search, LSH-based algorithms usually require a large number of random hash functions and may increase privacy risks [35].

Privacy-enhanced variants of LSH are recently proposed by combining LSH with cryptographic or information-theoretic protocols [27]. The privacy protection framework proposed in [28] generates a partial query instance by omitting certain bits in one or more sub-hash values to increase the ambiguity of query information for the server. The hash values of retrieved candidates are returned to the client for refinement. The framework is extended in [29] where partial encryption is performed on the hash code of each item to prevent an untrustworthy server from precisely linking queries and database records. In particular, the server uses the unencrypted part of each item for approximate indexing and search while the client uses the encrypted part for re-ranking the candidates received from the server. To limit the number of candidates sent to the client, the server performs a preliminary ranking based on the *partial distance* computed from the unencrypted part. The trade-off between privacy and utility of a search at the server side is therefore controlled by the number of unencrypted bits. We call this approach “LSH + partial distance” for brevity and will use it as one baseline approach for performance comparison.

3 PRIVACY-PRESERVING SIMILARITY SEARCH

In this section, we study the privacy-preserving similarity search problem in the framework of binary hypothesis testing. Section 3.1 details how we model the Hamming-ball

search problem as binary channel estimation. Section 3.2 derives the decision rule for the binary hypothesis test. Section 3.3 presents the design of the distance obfuscation mechanism and the obfuscated distance measure. Section 3.4 describes the test-based rank-ordered search scheme and discusses how to perform an approximate NN search based on the obfuscated distance measure.

3.1 Inference-Based Hamming-Ball Search

Consider a set Ω of binary strings, each of which represents a biometric identity and is of length D bits. Note that D is typically large for biometric data. The query \mathbf{q} is also of length D bits. As binary embedding is usually distance-preserving, we treat two biometric identities as similar if their Hamming distance is close.

Any string $\mathbf{p} \in \Omega$ is considered to be an r -neighbour of the query \mathbf{q} if the Hamming distance between \mathbf{p} and \mathbf{q} , denoted by $d(\mathbf{p}, \mathbf{q})$, satisfies $d(\mathbf{p}, \mathbf{q}) \leq r$. All strings $\mathbf{p} \in \Omega$ that satisfy $d(\mathbf{p}, \mathbf{q}) \leq r$ constitute the set, denoted by $\mathcal{B}(\mathbf{q}; r)$, of r -neighbours of the query \mathbf{q} within its neighbourhood of Hamming radius r . This is known as the r -neighbour detection or Hamming-ball search problem [15]. A related problem, known as k -NN search, is to find k records in Ω that are closest in Hamming distance to the query \mathbf{q} . We first resolve the r -neighbour detection problem, and later show how it can be extended to handle the k -NN search problem.

Our inference-based approach for r -neighbour detection requires a distance obfuscation mechanism to conceal the exact value of $d(\mathbf{p}, \mathbf{q})$ in a certain interval $[a, b]$, where both variables a and b take on integer values and satisfy $a \leq d(\mathbf{p}, \mathbf{q}) \leq b$. In this way, it is clear that $d(\mathbf{p}, \mathbf{q}) > r$ if $r < a$ and $d(\mathbf{p}, \mathbf{q}) < r$ if $r > b$. There is an uncertainty in comparing $d(\mathbf{p}, \mathbf{q})$ with r when $a \leq r \leq b$. In this case, we consider that an estimate of $d(\mathbf{p}, \mathbf{q})$, denoted by $\hat{d}(\mathbf{p}, \mathbf{q})$, is randomly chosen from the interval $[a, b]$ with equal probability. Then,

$$\Pr\{\hat{d}(\mathbf{p}, \mathbf{q}) \leq r\} = \frac{r - a + 1}{b - a + 1}. \quad (1)$$

For convenience, let

$$\pi \stackrel{\text{def}}{=} \frac{r - a + 1}{b - a + 1}, \quad (2)$$

from which we have $1/(b - a + 1) \leq \pi \leq 1$ for $a \leq r \leq b$, and the value of π depends on a and b for any given r .

We make inference based on the variable π . Specifically, let $x = 0$ if $d(\mathbf{p}, \mathbf{q}) > r$ and $x = 1$ if $d(\mathbf{p}, \mathbf{q}) \leq r$. Let Y be a binary random variable taking on values $y = 0$ if $\pi < \eta$ and $y = 1$ if $\pi \geq \eta$, where $\eta \in (0, 1)$ is a predefined threshold. We resolve the r -neighbour detection problem as a binary hypothesis test $H_0 : x = 0, H_1 : x = 1$ that predicts the original binary input x based on a decision rule δ as a function $\delta(y)$ of the values of Y .

This can be regarded as a binary channel estimation problem [36] where a binary digit x is transmitted over some noisy channel that outputs a random observation Y with outcome being either 0 or 1, as illustrated in Fig. 1. Due to uncertainties introduced by distance obfuscation and threshold comparison, the original binary input x may be received as $Y \neq x$ with probabilities $\lambda_0 = \Pr\{Y = 1|x = 0\}$




Fig. 1. Inference-based r -neighbour detection modelled as binary channel estimation with decision rule δ .

and $\lambda_1 = \Pr\{Y = 0|x = 1\}$. We wish to find an optimal way to decide what was transmitted based on the values of Y . In the following section, we design the decision rule δ given the “transmission” error probabilities λ_0 and λ_1 .

3.2 Design of the Decision Rule δ

Any decision rule δ is associated with two types of errors in testing the *null* hypothesis (i.e., H_0) versus the *alternative* hypothesis (i.e., H_1). Let $P_F(\delta)$ denote the false-alarm probability that H_1 is falsely accepted when H_0 is true. Let $P_M(\delta)$ denote the miss probability that H_1 is falsely rejected when H_1 is true. The detection probability $P_D(\delta) = 1 - P_M(\delta)$ is also called the *power* of the decision rule δ . There is a fundamental trade-off between $P_F(\delta)$ and $P_M(\delta)$ in the hypothesis test. According to the Neyman-Pearson criteria [37], the optimal decision rule is the one to the constrained optimization problem: $\max_{\delta} P_D(\delta)$ subject to $P_F(\delta) \leq \alpha$, with α being the *significance level* of the test. The Neyman-Pearson lemma [37] states that the likelihood ratio test is the most powerful test.

Consider the binary channel estimation problem in Fig. 1. The observation Y has probability mass functions $p_x(y) = \lambda_x$ if $y \neq x$ and $p_x(y) = 1 - \lambda_x$ if $y = x$, for $x = 0, 1$. The likelihood ratio for an observation $Y = y$ is thus given by

$$L(y) = \frac{p_1(y)}{p_0(y)} = \begin{cases} \frac{\lambda_1}{1-\lambda_0}, & \text{if } y = 0 \\ \frac{1-\lambda_1}{\lambda_0}, & \text{if } y = 1. \end{cases} \quad (3)$$

The Neyman-Pearson test rejects H_0 for $L(Y) > \eta$ for some η , and η is chosen such that $\Pr\{L(Y) > \eta\} = \alpha$ if H_0 is true. When $\lambda_0 + \lambda_1 < 1$, the Neyman-Pearson decision rule, which is the conditional probability of accepting H_1 given that we observe $Y = y$, turns out to be [36]

$$\tilde{\delta}_{\text{NP}}(y) = \begin{cases} \frac{\alpha}{\lambda_0}, & \text{if } y = 1 \\ 0, & \text{if } y = 0, \end{cases} \quad (4)$$

for $0 \leq \alpha < \lambda_0$, and

$$\tilde{\delta}_{\text{NP}}(y) = \begin{cases} 1, & \text{if } y = 1 \\ \frac{\alpha - \lambda_0}{1 - \lambda_0}, & \text{if } y = 0, \end{cases} \quad (5)$$

for $\lambda_0 \leq \alpha \leq 1$. The resulting detection probability is

$$P_D(\tilde{\delta}_{\text{NP}}) = \begin{cases} \alpha \left(\frac{1-\lambda_1}{\lambda_0} \right), & \text{if } 0 \leq \alpha < \lambda_0 \\ 1 - \lambda_1 + \lambda_1 \left(\frac{\alpha - \lambda_0}{1 - \lambda_0} \right), & \text{if } \lambda_0 \leq \alpha < 1. \end{cases} \quad (6)$$

It can be seen in (5) that $\tilde{\delta}_{\text{NP}}(y) = y$ when $\alpha = \lambda_0$. In other words, by allowing $P_F(\delta) \leq \lambda_0$, the simple decision rule $\delta(y) = y$ that accepts H_x if $y = x$ yields the same performance as the most powerful test according to the Neyman-

Pearson criteria. The resulting detection probability is $P_D(\delta) = 1 - \lambda_1$. In this case, the inherent transmission error probability serves as an upper bound of the hypothesis testing error probability, i.e., $P_F(\delta) + P_M(\delta) \leq \lambda_0 + \lambda_1$. This enables us to design the binary channel directly through a trade-off between the channel characteristics λ_0 and λ_1 .

3.3 Design of the Binary Channel

As discussed in Section 3.1, the binary channel in our context is in the form of a noisy channel that consists of distance obfuscation followed by threshold comparison. The former maps a distance value into an interval $[a, b]$ that defines π in (2). The latter yields λ_0 and λ_1 under the null and alternative hypotheses, respectively. Here, we present a distance obfuscation mechanism based on matching piecewise obfuscated binary codes in a carefully designed *multi-index multi-probe* (MIMP) scheme for hash-based similarity search.

Specifically, for all strings $\mathbf{p} \in \Omega$ and the query \mathbf{q} , we divide each of them into L non-overlapping segments, denoted by $\{\mathbf{p}^{(1)}, \mathbf{p}^{(2)}, \dots, \mathbf{p}^{(L)}\}$ and $\{\mathbf{q}^{(1)}, \mathbf{q}^{(2)}, \dots, \mathbf{q}^{(L)}\}$. By definition, we have $\mathbf{p} = \mathbf{p}^{(1)} \parallel \mathbf{p}^{(2)} \parallel \dots \parallel \mathbf{p}^{(L)}$ and $\mathbf{q} = \mathbf{q}^{(1)} \parallel \mathbf{q}^{(2)} \parallel \dots \parallel \mathbf{q}^{(L)}$, where \parallel is the concatenation operator. If the string length D is divisible by L , each substring is of length $s = D/L$ bits. If D is not divisible by L , each of the first $(D \bmod L)$ substrings is of length $s = \lceil D/L \rceil$ bits, and each of the remaining substrings has $s - 1$ bits.

For each substring $\mathbf{p}^{(i)}$ of \mathbf{p} , $i = 1, 2, \dots, L$, we create a set $\mathcal{V}_{\mathbf{p}^{(i)}}$ containing $\mathbf{p}^{(i)}$ and a one-bit variant of $\mathbf{p}^{(i)}$ that differs from $\mathbf{p}^{(i)}$ by one bit at a randomly selected position. For convenience, we write $\mathcal{V}_{\mathbf{p}^{(i)}} = \{\tilde{\mathbf{p}}_j^{(i)}\}$ where $j = 0, 1$, with $\tilde{\mathbf{p}}_0^{(i)} = \mathbf{p}^{(i)}$ and $\tilde{\mathbf{p}}_1^{(i)}$ being the one-bit variant of $\mathbf{p}^{(i)}$. For each substring $\mathbf{q}^{(i)}$ of the query \mathbf{q} , $i = 1, 2, \dots, L$, we create a set that contains all possible one-bit variants of $\mathbf{q}^{(i)}$, denoted by $\mathcal{V}_{\mathbf{q}^{(i)}} = \{\tilde{\mathbf{q}}_k^{(i)}\}$, where $k = 1, 2, \dots, s$ if the substring is of length s bits and $k = 1, 2, \dots, s - 1$ if the substring has $s - 1$ bits. We call a match between $\tilde{\mathbf{p}}_j^{(i)} \in \mathcal{V}_{\mathbf{p}^{(i)}}$ and $\tilde{\mathbf{q}}_k^{(i)} \in \mathcal{V}_{\mathbf{q}^{(i)}}$ as a *substring collision*. Let $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}})$ denote the *collision count* (i.e., the number of all possible substring collisions) between $\mathcal{V}_{\mathbf{p}^{(i)}}$ and $\mathcal{V}_{\mathbf{q}^{(i)}}$.

Proposition 1. *The collision count between $\mathcal{V}_{\mathbf{p}^{(i)}}$ and $\mathcal{V}_{\mathbf{q}^{(i)}}$ satisfies $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) \leq 1$ for all i . The case $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$ implies $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \leq 2$. The case $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 0$ implies $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \geq 2$.*

Proof. The following cases hold for each $i = 1, 2, \dots, L$.

- If $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) = 0$, we have $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$, since $\tilde{\mathbf{p}}_0^{(i)} \neq \tilde{\mathbf{q}}_k^{(i)}$ for all k and there exists exactly one $\tilde{\mathbf{q}}_k^{(i)} \in \mathcal{V}_{\mathbf{q}^{(i)}}$ such that $\tilde{\mathbf{p}}_1^{(i)} = \tilde{\mathbf{q}}_k^{(i)}$.
- If $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) = 1$, we have $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$, since $\tilde{\mathbf{p}}_1^{(i)} \neq \tilde{\mathbf{q}}_k^{(i)}$ for all k and there exists exactly one $\tilde{\mathbf{q}}_k^{(i)} \in \mathcal{V}_{\mathbf{q}^{(i)}}$ such that $\tilde{\mathbf{p}}_0^{(i)} = \tilde{\mathbf{q}}_k^{(i)}$.
- If $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) = 2$, we have $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) \leq 1$ where the equality holds if and only if the one-bit variant of $\mathbf{p}^{(i)}$ flips at a bit position at which $\mathbf{p}^{(i)}$ and $\mathbf{q}^{(i)}$ differ.
- If $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) > 2$, we have $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 0$, since $\tilde{\mathbf{p}}_j^{(i)} \neq \tilde{\mathbf{q}}_k^{(i)}$ for all j, k . \square

Corollary 1. *The sum of collision count $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}})$ over all i is at most L , i.e., $\sum_{i=1}^L C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) \leq L$.*

Corollary 2. *Since $0 \leq d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \leq s$, we have $0 \leq d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \leq 2$ given $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$, and $2 \leq d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \leq s$ given $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 0$.*

Proposition 2. *If $\mathbf{p} \in \mathcal{B}(\mathbf{q}; r)$ with $r < 2L$, then there exists at least one i , $1 \leq i \leq L$, such that $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$.*

Proof. Suppose that $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 0$ for all i . By Proposition 1, we must have $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \geq 2$ for all i . Then, $d(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^L d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \geq 2L$. That is, we have $\mathbf{p} \notin \mathcal{B}(\mathbf{q}; r)$ with $r < 2L$, which gives a contradiction. \square

Let m denote the sum of collision count $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}})$ over all i , i.e.,

$$m = \sum_{i=1}^L C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}). \quad (7)$$

Proposition 2 suggests that, under the condition $r < 2L$, the MIMP scheme ensures that any r -neighbor of the query \mathbf{q} has a non-zero value of m . We shall see that m serves as an obfuscated distance measure and is key to the design of the inference-based r -neighbor detection approach in our context. Note that $m \leq L$ from Corollary 1. Without loss of generality, let us assume that $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$ for $i = 1, 2, \dots, m$ and $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 0$ for $i = m+1, \dots, L$. Then, by Corollary 2, we have $0 \leq \sum_{i=1}^m d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \leq 2m$ and $2(L-m) \leq \sum_{i=m+1}^L d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) \leq s(L-m)$. Since by definition

$$d(\mathbf{p}, \mathbf{q}) = \sum_{i=1}^m d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}) + \sum_{i=m+1}^L d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)}), \quad (8)$$

we know that $d(\mathbf{p}, \mathbf{q})$ is within the range given by

$$2(L-m) \leq d(\mathbf{p}, \mathbf{q}) \leq s(L-m) + 2m. \quad (9)$$

Substituting $a = 2(L-m)$ and $b = s(L-m) + 2m$ in (2), we obtain

$$\pi = \frac{2 - (2L - 1 - r)/m}{(sL - 2L + 1)/m + 4 - s}, \quad (10)$$

for $2(L-m) \leq r \leq s(L-m) + 2m$ or, equivalently,

$$L - \frac{r}{2} \leq m \leq \frac{sL - r}{s - 2}. \quad (11)$$

Note that, with $r < 2L$, we must have $(sL - r)/(s - 2) > L$. Since $m \leq L$ by Corollary 1, we rewrite (11) as

$$L - \frac{r}{2} \leq m \leq L. \quad (12)$$

In addition, with $s \geq 2$, π in the form of (10) monotonically increases as m increases. Accordingly, based on the definition of π , a larger value of m corresponds to a higher probability that \mathbf{p} is an r -neighbour of \mathbf{q} .

As a result of (12), any string $\mathbf{p} \in \Omega$ with $m < L - r/2$ is not an r -neighbour of the query \mathbf{q} . On the other hand, all strings $\mathbf{p} \in \Omega$ with m being in the inference region defined




Fig. 2. Empirical probability mass functions of π under H_0 and H_1 , respectively, with the UBIRIS dataset, given $L = 30$ and $r = 50$. Choosing $\eta = 0.1144$ minimizes $\lambda_0 + \lambda_1$.

by (12) are subjected to the hypothesis testing, from which we can obtain through (10) the empirical probability mass functions of π under hypothesis H_0 and hypothesis H_1 , respectively, and then choose a value for the threshold η to set the channel characteristics λ_0 and λ_1 .

For example, let us consider the UBIRIS dataset [38] and use 400 eyes each with 10 iris images. Each iris image is represented by an iris code of 400 bits generated using a rotation invariant ordinal feature coding scheme [13]. We enroll each eye with one iris code at a time, and then use another iris code from the same eye to search. This is run for 24 times with different enrollments and queries, which results in a total of 3,840,000 pairwise comparisons, including 6,026 neighbouring pairs and 3,833,974 non-neighbouring pairs as ground truth, given the Hamming radius $r = 50$.

We set $L = 30$ and perform the MIMP scheme in each run. All enrolled strings whose m value with respect to corresponding queries satisfying (12), which is $5 \leq m \leq 30$ in this example, are subjected to inference. This involves 6,026 neighbouring pairs and 5,859 non-neighbouring pairs. Comparing with the ground truth, the inference set includes all neighbouring pairs and a small fraction (i.e., 0.15 percent) of non-neighbouring pairs. By (10), the π values can be evaluated for these instances. Fig. 2 plots the empirical probability mass functions of π under H_0 and H_1 , respectively.

We choose the threshold η to minimize $\lambda_0 + \lambda_1$ for the binary channel, which we recall is the upper bound of the hypothesis testing error probability $P_F(\delta) + P_M(\delta)$ as a result of the chosen decision rule $\delta(y) = y$ discussed in Section 3.2. Note that we may also choose a smaller η value to increase the detection probability $P_D(\delta)$, i.e., $1 - \lambda_1$, by allowing a larger λ_0 . In Fig. 2, $\eta = 0.1144$, which yields $\lambda_0 = 0.082$ and $\lambda_1 = 0.036$. Given the η value, the inference-based approach determines if an enrolment is an r -neighbour of the query by testing if its π value satisfies $\pi \geq \eta$. Accordingly, taking into account all the 3,840,000 pairwise comparisons in this example, the overall detection probability is 96.4 percent and the overall false-alarm probability is 0.0125 percent.

3.4 Test-Based Rank-Ordered Search

Conventionally, a rank-ordered search is done by ordering the retrieved candidates according to their distance values. In this section, we present a test-based rank-ordered search scheme without the need of comparing distance values. The basic idea is to apply our proposed inference-based r -neighbour detection approach and evaluate a series of nested inference regions with sequentially increased values of the Hamming radius r .

TABLE 1
Inference-Based r -Neighbour Detection
over the UBIRIS Dataset

Hamming radius r	30	40	50
Inference region \mathcal{I}	[15, 30]	[10, 30]	[5, 30]
Threshold η for π	0.1064	0.1110	0.1144
Threshold μ for m	22	19	16
Overall detection probability	96.5%	95.9%	96.4%
Overall false-alarm probability	0.0186%	0.0140%	0.0125%

Take the UBIRIS dataset for example. Let us consider three values for the Hamming radius r as shown in Table 1, i.e., $r_1 = 30$, $r_2 = 40$, $r_3 = 50$. We perform the MIMP scheme with $L = 30$. Thus, the inference region \mathcal{I} defined by (12) for each corresponding r value is $\mathcal{I}_1 = [15, 30]$, $\mathcal{I}_2 = [10, 30]$, $\mathcal{I}_3 = [5, 30]$, respectively. Note that in the case of the MIMP scheme, as long as L is fixed, the m value of any enrollment remains the same regardless of the r value. For each value of r , we obtain through (10) the empirical probability mass functions of π , and choose the threshold η to minimize $\lambda_0 + \lambda_1$. For the UBIRIS dataset, it turns out that $\eta_1 = 0.1064$, $\eta_2 = 0.1110$, $\eta_3 = 0.1144$, respectively.

Once the η values are learned, we can determine if an enrollment \mathbf{p} is an r -neighbour of the query \mathbf{q} by testing if the π value of \mathbf{p} satisfies $\pi \geq \eta$. Given π in the form of (10), it can be shown that testing $\pi \geq \eta$ is equivalent to testing

$$m \geq \frac{\eta(sL - 2L + 1) + 2L - 1 - r}{2 + \eta(s - 4)} \stackrel{\text{def}}{=} \mu. \quad (13)$$

Note in (13) that it suffices for μ to assume integer values since m in this context takes on integer values only. Thus, for the UBIRIS dataset, we have $\mu_1 = 22$, $\mu_2 = 19$, $\mu_3 = 16$, respectively. The resulting detection probability and false-alarm probability of r -neighbour detection are provided in Table 1 for each value of r .

Given $\mu_1 > \mu_2 > \mu_3$ and since the inference region \mathcal{I} is nested, i.e., $\mathcal{I}_1 \subset \mathcal{I}_2 \subset \mathcal{I}_3$, the inferred set $\mathcal{B}'(\mathbf{q}; r)$ of r -neighbours of the query \mathbf{q} in this case must also be nested, i.e., $\mathcal{B}'(\mathbf{q}; r_1) \subset \mathcal{B}'(\mathbf{q}; r_2) \subset \mathcal{B}'(\mathbf{q}; r_3)$. In general, if we consider G sequentially increased values of the Hamming radius r and given that the inferred set of r -neighbours is nested, we can order the retrieved candidates in G ranks. In particular, those in the set $\mathcal{B}'(\mathbf{q}; r_1)$ are designated as Rank 1 candidates. Subsequently, for $g = 2, 3, \dots, G$, those in $\mathcal{B}'(\mathbf{q}; r_g) - \mathcal{B}'(\mathbf{q}; r_{g-1})$ form the set of Rank g candidates.

As a special case, an approximate k -NN search can be performed by finding the smallest G such that the set $\mathcal{B}'(\mathbf{q}; r_G)$ contains at least k candidates. We return all candidates in the top $G - 1$ ranks, and the remaining candidates are chosen from Rank G candidates arbitrarily. The latter is due to the nature of binary hypothesis testing where candidates in the same rank are considered to be of the same similarity level to the query regardless of their m values. It is interesting to observe that, in the case of the test-based rank-ordered search, given $\mu_i > \mu_j$, the m values of Rank i candidates must be larger than those of Rank j candidates. This observation motivates us to further consider a simple implementation of k -NN search where we simply choose

the top k candidates in descending order of the m value without the need of test-based ranking.

4 PRIVACY-PRESERVING INDEX GENERATION

In Section 3, we introduced an inference-based approach for similarity search. The MIMP scheme works by matching piecewise obfuscated binary codes and enables hash-based indexing of biometric identities. This can be regarded as *substring obfuscation* by introducing controlled “noise” to piecewise biometric data before disseminating them over hash tables. However, as shown in Section 5.3, from the perspective of privacy protection, such segments of biometric data should not be used directly as search indexes and stored in hash tables. Cryptographic hashing may be used to generate secure signatures. However, standard cipher codes typically have hundreds of bits, e.g., 256 bits by SHA-256, while codes of length longer than 32 bits are not suitable for indexing data structures [15]. Subdivision of cipher codes does not help as Hamming distances are not preserved after encryption. In this paper, we address the problem by indexing in randomized Montgomery domains.

4.1 Indexing in Montgomery Domains

Montgomery multiplication has been used in cryptographic schemes mainly for implementing fast modulo operations with large-integer arithmetic [39]. Here, we exploit its elementary form to generate randomized signatures. Specifically, let N be a prime number, and let R be a positive integer that is coprime to N , i.e., $\gcd(R, N) = 1$. The *Montgomery form* of an integer x is defined as

$$M(x; R, N) = xR \bmod N, \quad (14)$$

where N is the modulus and R is the multiplier. We also call $M(x; R, N)$ as the (R, N) -residue of x following Definition 9.2.2 in [39]. In our context, for a binary string $\mathbf{p}^{(i)}$ of length s bits, its Montgomery form can be evaluated by the sum of (R, N) -residues at each one-bit position in $\mathbf{p}^{(i)}$. That is,

$$\begin{aligned} M(\mathbf{p}^{(i)}; R, N) &= \left(\sum_{b=1}^s \mathbf{p}^{(i)}[b] \cdot 2^{b-1} \right) R \bmod N \\ &= \sum_{b=1}^s \mathbf{p}^{(i)}[b] \cdot M(2^{b-1}; R, N) \bmod N, \end{aligned} \quad (15)$$

where $\mathbf{p}^{(i)}[b] \in \{0, 1\}$. The reference $M(2^{b-1}; R, N)$ may be pre-computed for large-integer arithmetic. We regard (15) as a projection of $\mathbf{p}^{(i)}$ into the Montgomery domain of (R, N) -residues, also called the (R, N) -domain for brevity. Note that Montgomery multiplication does not preserve the original distance values and thus the resulting signatures can withstand adversarial similarity analysis.

We consider a three-party scenario that involves a user, a data owner and a server, which is typical in a networked computing environment like cloud [21]. The data owner is responsible for user enrolment, server registration and index generation. The server maintains the hash tables and provides the computation-intensive task of similarity search on behalf of the data owner. Upon a query from the user, the server retrieves the most likely candidates and returns the search results to the user. Fig. 3 illustrates the main

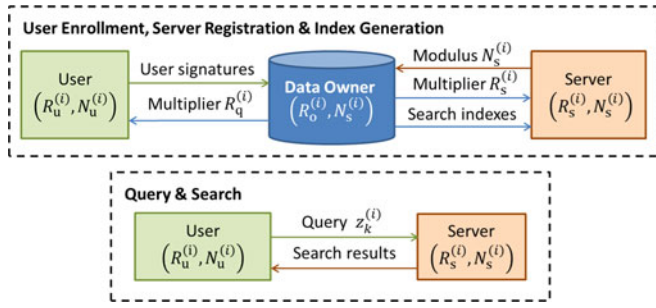


Fig. 3. Main procedures of the three-party search protocol.

procedures of our three-party protocol that enables similarity search in randomized Montgomery domains. They are described below in more details.

User Enrolment. The user generates the binary string \mathbf{p} from biometric data and divides it into L substrings $\{\mathbf{p}^{(i)}\}$, each of which is associated with a set $\mathcal{V}_{\mathbf{p}^{(i)}}$ as described in Section 3.3. For each chosen modulus $N_u^{(i)}$, the user randomly generates a multiplier $R_u^{(i)}$ that is coprime to $N_u^{(i)}$, and maps all $\tilde{\mathbf{p}}_j^{(i)} \in \mathcal{V}_{\mathbf{p}^{(i)}}$ into the $(R_u^{(i)}, N_u^{(i)})$ -domain. Then, the user sends $M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)})$ as the signature of $\tilde{\mathbf{p}}_j^{(i)}$ for all i, j to the data owner.

Server Registration. For provision of service, the server registers at the data owner with a set of moduli $\{N_s^{(i)}\}$. Then, for each i , the data owner randomly generates a multiplier $R_o^{(i)}$ that is coprime to $N_s^{(i)}$ for index generation.

Index Generation. The data owner maps the user's signature $M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)})$ into the $(R_o^{(i)}, N_s^{(i)})$ -domain, and then sends the nested Montgomery form $M(M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)}); R_o^{(i)}, N_s^{(i)})$ for all $\tilde{\mathbf{p}}_j^{(i)} \in \mathcal{V}_{\mathbf{p}^{(i)}}$ to the server for indexing $\mathbf{p}^{(i)}$ in the i th hash table. To enable the user to search at the server, the data owner randomly generates for each i a multiplier $R_q^{(i)}$ that is coprime to $N_s^{(i)}$, and sends $R_q^{(i)}$ for all i to the user. Let $R_q^{(i)}$ be a modular multiplicative inverse of $R_q^{(i)}$ modulo $N_s^{(i)}$, i.e.,

$$R_q^{(i)} R_q^{(i)} \bmod N_s^{(i)} = 1. \quad (16)$$

The data owner also computes

$$R_s^{(i)} = M(R_q^{(i)}; R_o^{(i)}, N_s^{(i)}), \quad (17)$$

and sends the resulting multiplier $R_s^{(i)}$ for all i to the server.

Query and Search. The user divides the query \mathbf{q} into L substrings and, for each substring $\mathbf{q}^{(i)}$, generates a set $\mathcal{V}_{\mathbf{q}^{(i)}}$ as described in Section 3.3. For each i , the user maps all $\tilde{\mathbf{q}}_k^{(i)} \in \mathcal{V}_{\mathbf{q}^{(i)}}$ into the $(R_u^{(i)}, N_u^{(i)})$ -domain, and computes

$$z_k^{(i)} = M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}) R_q^{(i)}, \quad (18)$$

using the multiplier $R_q^{(i)}$. The user sends $z_k^{(i)}$ to the server who in turn maps it into the $(R_s^{(i)}, N_s^{(i)})$ -domain. Note that the result $M(z_k^{(i)}; R_s^{(i)}, N_s^{(i)})$ and the nested Montgomery form of $\tilde{\mathbf{q}}_k^{(i)}$, i.e., $M(M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}); R_o^{(i)}, N_s^{(i)})$, are equivalent. This is because, by (15), (16), (17), and (18) and using the congruence relation $a(b \bmod n) \equiv ab \bmod n \pmod{n}$ that holds for all positive integers a, b, n , we have

$$\begin{aligned} M(z_k^{(i)}; R_s^{(i)}, N_s^{(i)}) &= M(M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}) R_q^{(i)}; R_s^{(i)}, N_s^{(i)}) \\ &= M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}) R_q^{(i)} R_q^{(i)} R_o^{(i)} \bmod N_s^{(i)} \bmod N_s^{(i)} \\ &= M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}) R_o^{(i)} (R_q^{(i)} R_q^{(i)} R_o^{(i)} \bmod N_s^{(i)}) \bmod N_s^{(i)} \\ &= M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}) R_o^{(i)} \bmod N_s^{(i)} \\ &= M(M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}); R_o^{(i)}, N_s^{(i)}). \end{aligned} \quad (19)$$

Thus, the server can use the $(R_s^{(i)}, N_s^{(i)})$ -residue of $z_k^{(i)}$ to probe the search indexes in the i th hash table and perform the MIMP scheme as described in Section 3.3.

4.2 Substring Collision in Montgomery Domains

Note that, given a substring collision between $\tilde{\mathbf{p}}_j^{(i)} \in \mathcal{V}_{\mathbf{p}^{(i)}}$ and $\tilde{\mathbf{q}}_k^{(i)} \in \mathcal{V}_{\mathbf{q}^{(i)}}$, we must have

$$\begin{aligned} M(M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)}); R_o^{(i)}, N_s^{(i)}) \\ = M(M(\tilde{\mathbf{q}}_k^{(i)}; R_u^{(i)}, N_u^{(i)}); R_o^{(i)}, N_s^{(i)}). \end{aligned} \quad (20)$$

However, having (20) does not necessarily imply $\tilde{\mathbf{p}}_j^{(i)} = \tilde{\mathbf{q}}_k^{(i)}$. Therefore, collision detection based on (20) may result in false positives but *no* false negatives. In other words, matching in Montgomery domains may increase the false-alarm probability but have *no* effect on the detection probability in our inference-based similarity search scheme. To reduce the false positives, we enhance the search protocol by indexing in Montgomery domains using *multiple* independently generated signatures rather than one. The idea can be conveniently explained by considering two arbitrary substrings of length s bits as follows.

Let x and y be the natural number representation of the two substrings, respectively. We first consider T randomly chosen pairs of (R_t, N_t) , $t = 1, 2, \dots, T$, each yielding a Montgomery form for x and y . For convenience, let $\gamma_t(x) = M(x; R_t, N_t)$ and $\gamma_t(y) = M(y; R_t, N_t)$, and let $\gamma(x) = \{\gamma_t(x)\}$ and $\gamma(y) = \{\gamma_t(y)\}$ denote the set of T randomized signatures for x and y , respectively. In this way, we consider the two substrings matched, i.e., $x = y$, if and only if $\gamma(x) = \gamma(y)$, i.e., $\gamma_t(x) = \gamma_t(y)$ for all t .

Suppose that R_t is encoded by c_R bits, i.e., $0 < R_t < 2^{c_R}$, and N_t is encoded by c_N bits, i.e., $0 < N_t < 2^{c_N}$. Then, the binary code of xR_t and yR_t has $s + c_R$ bits. Let $\nu(2^{c_N})$ be the number of primes in the range $(0, 2^{c_N})$. We choose each modulus N_t independently from the $\nu(2^{c_N})$ prime numbers at random with equal probability. A prime number is “bad” for matching x and y if the random choice of N_t results in a false positive, i.e., $\gamma_t(x) = \gamma_t(y)$ given $x \neq y$. It is known that, for $s + c_R < \nu(2^{c_N})$, there can be at most $s + c_R - 1$ such “bad” prime numbers in the range $(0, 2^{c_N})$ [40]. Since every prime number has an equal probability of being chosen, the probability $\Pr\{\gamma_t(x) = \gamma_t(y) | x \neq y\}$ is bounded by

$$\frac{s + c_R - 1}{\nu(2^{c_N})} \stackrel{\text{def}}{=} \beta(s; c_R, c_N). \quad (21)$$

The false positive probability $\Pr\{\gamma(x) = \gamma(y) | x \neq y\}$ is therefore bounded by $\beta(s; c_R, c_N)^T$, which can be controlled by tuning the parameters c_R, c_N and T for a given s .




Fig. 4. Bound $\beta(s + c_N; 2c_R, c_N)^{T^2}$ with respect to the modulus encoding length c_N for $T = 1, 2, 3$, given $s = 14$ and $c_R = 15$.

Next, consider the nested Montgomery form for x and for y . That is, for each $\gamma_t(x)$ and $\gamma_t(y)$, we further choose T pairs of (R_v, N_v) , $v = 1, 2, \dots, T$, independently at random. Each pair of (R_v, N_v) yields the nested Montgomery form $M(\gamma_t(x); R_v, N_v)$ for x and $M(\gamma_t(y); R_v, N_v)$ for y . For convenience, let $\psi_{tv}(x) = M(\gamma_t(x); R_v, N_v)$ and $\psi_{tv}(y) = M(\gamma_t(y); R_v, N_v)$. Accordingly, we generate a set of T^2 randomized signatures for x and y , denoted by $\Psi(x) = \{\psi_{tv}(x)\}$ and $\Psi(y) = \{\psi_{tv}(y)\}$, respectively. In this way, we consider $x = y$ if and only if $\Psi(x) = \Psi(y)$, i.e., $\psi_{tv}(x) = \psi_{tv}(y)$ for all t and v . Since $N_t < 2^{c_N}$, both $\gamma_t(x)$ and $\gamma_t(y)$ can be encoded by c_N bits. Setting $s = c_N$ in (21) yields $\Pr\{\psi_{tv}(x) = \psi_{tv}(y) | \gamma_t(x) \neq \gamma_t(y)\} \leq \beta(c_N; c_R, c_N)$. Then, it can be shown that $\Pr\{\psi_{tv}(x) = \psi_{tv}(y) | x \neq y\}$ is bounded by $\beta(s + c_N; 2c_R, c_N)$. The resulting false positive probability $\Pr\{\Psi(x) = \Psi(y) | x \neq y\}$ in the case of using nested Montgomery forms as search indexes is therefore bounded by $\beta(s + c_N; 2c_R, c_N)^{T^2}$.

Fig. 4 plots the bound $\beta(s + c_N; 2c_R, c_N)^{T^2}$ with respect to the parameter c_N for $T = 1, 2, 3$, given $s = 14$ and $c_R = 15$. It can be seen that the bound reduces drastically as c_N or T increases. For example, with $T = 2$, $c_N = 15$ and hence $v(2^{15}) = 3512$, the false positive probability is less than 7.44×10^{-8} . In this paper, we use $T = 2$ and $c_N = 15$ in our experimental settings.

5 PRIVACY-PRESERVING STRENGTH

In this section, we study the privacy-preserving strength of our inference-based framework for similarity search in randomized Montgomery domains. In particular, we are concerned with adversarial learning of the biometric similarity information that can be gleaned from the search indexes and the retrieval process. We begin by presenting in Section 5.1 a security analysis of the three-party search protocol to understand what an attacker can do if they know the various primitives (i.e., multipliers and moduli) in each indexing step. We argue that the privacy threats may all boil down to the problem of how hard it is for the attacker to infer the biometric identity from the user's signatures in Montgomery forms. This motivates us to apply an information-theoretic approach for evaluating in Section 5.2 the information leakage in Montgomery domains and in Section 5.3 the privacy gain in index generation. We further provide in Section 5.4 a discussion on the well-known trade-off between privacy and utility that is applied also in our context of similarity search.

5.1 Security Analysis of the Search Protocol

In our design of the three-party search protocol, we consider a model where the different parties do not necessarily trust each other. We also assume that 1) the three parties do not collude with each other, 2) the communication channel between the parties is secure (e.g., encrypted), and 3) each party may behave in a curious-but-honest way, i.e., it follows the search protocol but may use any data in possession to glean additional information [5], [6], [28], [29]. As shown in Fig. 3, the three parties operate in separate Montgomery domains defined by the corresponding primitive pair, i.e., $(R_u^{(i)}, N_u^{(i)})$ for the user, $(R_o^{(i)}, N_s^{(i)})$ for the data owner, and $(R_s^{(i)}, N_s^{(i)})$ for the server. The following discusses the privacy threat in each main procedure of the search protocol.

User Enrolment. The attacker may seize user signatures in the form of $M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)})$, but is not likely to know the user's primitive pair $(R_u^{(i)}, N_u^{(i)})$, as the latter is kept private to the user. In Section 5.2, we show that recovering $\tilde{\mathbf{p}}_j^{(i)}$ from its Montgomery form without knowing the primitives can be made almost as difficult as a wild guess. On the other hand, even if the attacker is able to crack user signatures, it needs to do so for a sufficiently large number of users in order to derive the underlying biometric similarity distribution. This makes the inference cost of adversarial learning even more prohibitive.

Server Registration. We note that, if different application servers use the same set of moduli $\{N_s^{(i)}\}$ to register provision of service at the data owner, the attacker can track users covertly by cross-matching over different applications, known as *linkage attack* [4]. To deal with this privacy threat, the servers can choose different sets of moduli $\{N_s^{(i)}\}$, so that the data owner can produce distinct sets of search indexes for the same biometric database.

Index Generation. The attacker may sniff all the search indexes generated by the data owner. Since the search indexes are generated in the nested Montgomery form $M(M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)}); R_o^{(i)}, N_s^{(i)})$, the attacker cannot derive the biometric similarity distribution directly from the search indexes, as Montgomery multiplication does not preserve the original distance values. The attacker may alternatively attempt to glean the information by inferring each biometric identity from the search indexes and then performing pairwise comparisons for distance computation. In this way, the attacker needs to compromise two layers of Montgomery multiplications defined in the respective Montgomery domains. We note that secure communications can prevent the attacker from obtaining the primitives $N_s^{(i)}$, $R_s^{(i)}$ and $R_q^{(i)}$. Otherwise, the attacker may derive the data owner's primitive $R_o^{(i)}$, which will enable the attacker to obtain $M(\tilde{\mathbf{p}}_j^{(i)}; R_u^{(i)}, N_u^{(i)})$. The privacy threat is again reduced to cracking the user's signatures as discussed before.

Query and Search. Recall that the obfuscated distance measure is key to our design of the inference-based similarity search scheme. It enables the server to retrieve most likely candidates without knowing the exact distance values. The attacker may attempt to glean the similarity information from the obfuscated distance measure. As a result, the risk of data disclosure in the retrieval process also depends on how different the distribution of the obfuscated




Fig. 5. $I(X; \Gamma_1, \Gamma_2)$ with respect to the multiplier encoding length c_R for $s = 4, 7, 14$. The mean and standard deviation are obtained with ten pairs of prime moduli N_1 and N_2 randomly chosen in $(0, 2^{15})$.

distance values can be made from that of the exact distance values, which we will discuss in Section 5.4.

5.2 Information Leakage in Montgomery Domains

In privacy studies, *mutual information* is often used to measure the average risk of data disclosure [41]. Accordingly, we quantify information leakage in our context as the mutual information between a substring data variable and its independently generated signatures in Montgomery domains. Specifically, consider a substring of length s bits whose natural number representation is a discrete random variable X with probability mass function $p(x) = 2^{-s}$ for all x in the range $\mathcal{X} = [0, 2^s)$. The entropy of X , denoted by $H(X)$, is thus s bits [42]. For $t = 1, 2$ and given the modulus N_t , the Montgomery form of X , denoted by Γ_t , is also a discrete random variable. In particular, let $\mathfrak{R}(N_t)$ be the set of all eligible values of the multiplier R_t in the range $(0, 2^{c_R})$ that are coprime to N_t . Then, for $X = x$, every $R_t \in \mathfrak{R}(N_t)$ yields an (R_t, N_t) -residue, i.e., $\gamma_t = M(x; R_t, N_t)$. Let all possible values of Γ_t constitute the subset $\mathcal{G}_t \subset [0, N_t)$. The mutual information between X and its Montgomery forms Γ_1, Γ_2 , denoted by $I(X; \Gamma_1, \Gamma_2)$, is then given by [42]

$$I(X; \Gamma_1, \Gamma_2) = H(X) - H(X | \Gamma_1, \Gamma_2). \quad (22)$$

The conditional entropy $H(X | \Gamma_1, \Gamma_2)$ can be obtained by the *chain rule for entropy* [42] as

$$H(X | \Gamma_1, \Gamma_2) = H(\Gamma_1, \Gamma_2, X) - H(\Gamma_1, \Gamma_2), \quad (23)$$

where $H(\Gamma_1, \Gamma_2, X)$ and $H(\Gamma_1, \Gamma_2)$ can be obtained from the joint distribution $p(x, \gamma_1, \gamma_2)$ as follows. Since Γ_1 and Γ_2 are independent, we have

$$p(x, \gamma_1, \gamma_2) = p(\gamma_1, \gamma_2 | x) p(x) = p(x) \prod_{t=1}^2 p(\gamma_t | x), \quad (24)$$

where the conditional distribution $p(\gamma_t | x)$ can be derived by enumerating all $R_t \in \mathfrak{R}(N_t)$ with respect to N_t . The marginal distribution $p(\gamma_1, \gamma_2)$ can be obtained as

$$p(\gamma_1, \gamma_2) = \sum_{x \in \mathcal{X}} p(x, \gamma_1, \gamma_2). \quad (25)$$

Then, the joint entropy $H(\Gamma_1, \Gamma_2, X)$ is given by

$$H(\Gamma_1, \Gamma_2, X) = - \sum_{x \in \mathcal{X}} \sum_{\gamma_1 \in \mathcal{G}_1} \sum_{\gamma_2 \in \mathcal{G}_2} p(x, \gamma_1, \gamma_2) \log p(x, \gamma_1, \gamma_2), \quad (26)$$




Fig. 6. $I(X; \Gamma_1, \Gamma_2)$ with respect to the prime moduli N_1 and N_2 chosen to span the range $(0, 2^{15})$, given $s = 7$ and $c_R = 15$.

and the joint entropy $H(\Gamma_1, \Gamma_2)$ is given by

$$H(\Gamma_1, \Gamma_2) = - \sum_{\gamma_1 \in \mathcal{G}_1} \sum_{\gamma_2 \in \mathcal{G}_2} p(\gamma_1, \gamma_2) \log p(\gamma_1, \gamma_2). \quad (27)$$

Fig. 5 plots $I(X; \Gamma_1, \Gamma_2)$ with respect to the parameter c_R for $s = 4, 7, 14$. With $c_N = 15$, we randomly select ten pairs of moduli N_1 and N_2 from the $\nu(2^{15}) = 3512$ available primes. For each pair of N_1 and N_2 , and for each particular value of c_R and s , we calculate $I(X; \Gamma_1, \Gamma_2)$ and present in Fig. 5 the mean and standard deviation of $I(X; \Gamma_1, \Gamma_2)$. It can be seen that, as c_R increases, $I(X; \Gamma_1, \Gamma_2)$ monotonically decreases and approaches zero. The larger the value of s , the smaller is $I(X; \Gamma_1, \Gamma_2)$ when it converges.

Fig. 6 shows the effect of modulus on $I(X; \Gamma_1, \Gamma_2)$. Given $s = 7$ and $c_R = 15$, the plot displays $I(X; \Gamma_1, \Gamma_2)$ with respect to the moduli N_1 and N_2 that are chosen to span the range $(0, 2^{15})$. In general, as N_1 or N_2 increases, $I(X; \Gamma_1, \Gamma_2)$ decreases and approaches zero. When both N_1 and N_2 are larger than 139, corresponding to 99 percent of prime numbers in the range $(0, 2^{15})$, $I(X; \Gamma_1, \Gamma_2)$ quickly drops to 0.066 bits. This indicates that information leakage in Montgomery domains can also be made negligibly small by choosing sufficiently large moduli values.

5.3 Privacy Gain in Index Generation

To measure the effect of applying privacy-preserving mechanisms in index generation, we introduce the notion of *privacy gain*, defined in our context as the increase in the uncertainty of inferring a biometric identity \mathbf{p} from the search indexes. Let $\hat{\mathbf{p}}$ be a random variable whose possible values are estimates of \mathbf{p} , which we recall is in the form of a binary string of length D bits. In an information-theoretic approach, the privacy gain can be measured by the entropy of $\hat{\mathbf{p}}$ conditional on the knowledge of the search indexes. Accordingly, the privacy gain is *null* if one directly uses the segments $\{\mathbf{p}^{(1)}, \mathbf{p}^{(2)}, \dots, \mathbf{p}^{(L)}\}$ for indexing \mathbf{p} as the adversary can recover $\mathbf{p} = \mathbf{p}^{(1)} || \mathbf{p}^{(2)} || \dots || \mathbf{p}^{(L)}$ unambiguously. On the other hand, the privacy gain is *maximum*, i.e., D bits, if the search indexes are generated in an uninformative way.

Let us consider what if we use exactly the original form of the obfuscated segments $\{\mathcal{V}_{\mathbf{p}^{(1)}}, \mathcal{V}_{\mathbf{p}^{(2)}}, \dots, \mathcal{V}_{\mathbf{p}^{(L)}}\}$ for hash-based indexing. It is clear in this case that, for any i , an estimate of $\mathbf{p}^{(i)}$ has only two outcomes in the set $\mathcal{V}_{\mathbf{p}^{(i)}}$, i.e., $\hat{\mathbf{p}}^{(i)} = \tilde{\mathbf{p}}_j^{(i)}$ for $j = 0, 1$ with equal probability. Thus, the conditional entropy $H(\hat{\mathbf{p}}^{(i)} | \mathcal{V}_{\mathbf{p}^{(i)}})$ is exactly one bit for all i , and hence the privacy gain in this case is L bits.

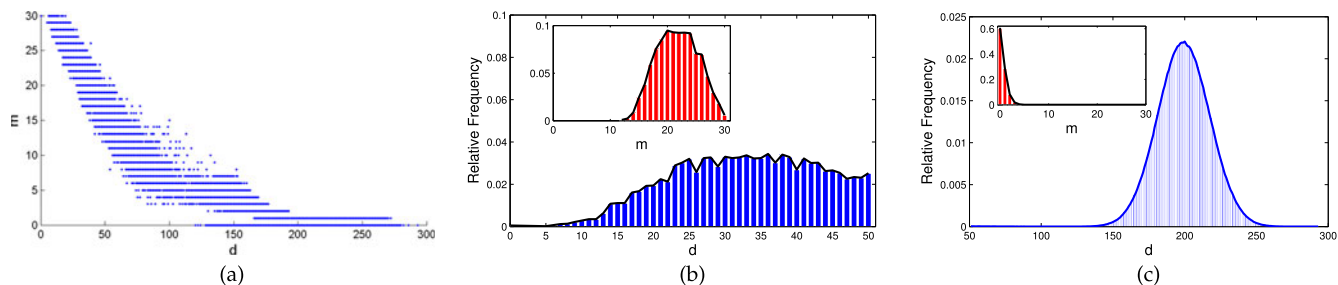


Fig. 7. Effect of distance obfuscation with $L = 30$ for the UBIRIS dataset. (a) Mapping from Hamming distance d to obfuscated distance measure m . (b) Histogram of d and that of m for neighbouring pairs with $d \leq 50$. (c) Histogram of d and that of m for non-neighbouring pairs with $d > 50$.

We recall that the baseline “LSH + partial distance” approach [29] generates the search indexes by performing random sampling of bit positions. As a result, the more the distinct positions that are sampled, the less is the uncertainty left in inferring the biometric identity. The privacy gain of the LSH-based approach is thus equal to the number of unknown bits. For comparison, we randomly sample an input string with $s = 8$ bits at a time for $L = 50$ times with replacement. We perform simulations over 100,000 input strings each of length 400 bits. The input bits are independently generated and each bit is 0 or 1 with equal probability. The results show that the privacy gain in such a context is 147 bits on average.

Now, consider our proposed index generation scheme using the Montgomery form of each substring $\tilde{\mathbf{p}}_j^{(i)} \in \mathcal{V}_{\mathbf{p}^{(i)}}$. In this context, if every $\tilde{\mathbf{p}}_j^{(i)}$ has two independently generated signatures in Montgomery domains, we have a total of four randomized signatures for indexing $\mathbf{p}^{(i)}$, denoted by $\Gamma_1^{(i)}, \Gamma_2^{(i)}, \Gamma_3^{(i)}, \Gamma_4^{(i)}$. In this case, let us again treat the natural number representation of $\tilde{\mathbf{p}}^{(i)}$ as a discrete random variable X with probability mass function $p(x) = 2^{-s}$ for all x in the range $\mathcal{X} = [0, 2^s)$. Then, in a similar manner to the way we calculate $H(X|\Gamma_1, \Gamma_2)$ in Section 5.2 through (23), (24), (25), (26), and (27), and setting $s = 8$, $c_R = 15$, $c_N = 15$, we obtain the mean value of the conditional entropy $H(\tilde{\mathbf{p}}^{(i)} | \Gamma_1^{(i)}, \Gamma_2^{(i)}, \Gamma_3^{(i)}, \Gamma_4^{(i)})$ as 7.82 bits and therefore the privacy gain is 391 bits for $L = 50$, which is rather close to $D = 400$ bits in this case.

5.4 Privacy-Utility Trade-Off in Similarity Search

As discussed in Section 3, the distance obfuscation mechanism in the MIMP scheme enables inference-based similarity search to make judicious test decisions based on the obfuscated distance measure without the need of directly evaluating the Hamming distance. For any particular pair of \mathbf{p} and \mathbf{q} , given the number of substrings L and the value of m that is obtained from (7), the Hamming distance $d(\mathbf{p}, \mathbf{q})$ is concealed in an interval given by (9) that depends on the value of m . This is analogous to the concept of data anonymization [31].

It is important to note that the parameter L plays a critical role in the trade-off between privacy and utility (i.e., search accuracy) in this context. We observe that:

- In the extreme case where $L = D$, each substring has exactly one bit. Since $d(\mathbf{p}^{(i)}, \mathbf{q}^{(i)})$ in this case is either 0 or 1, by Proposition 1, we must have $C(\mathcal{V}_{\mathbf{p}^{(i)}}, \mathcal{V}_{\mathbf{q}^{(i)}}) = 1$ for all i and hence $m = D$ regardless of the value of $d(\mathbf{p}, \mathbf{q})$. This case yields the

maximum privacy achievable as all strings $\mathbf{p} \in \Omega$ have the same m value and thus cannot be ranked. However, in this extreme case, the inference-based approach does not work as similarity search is reduced to a wild guess, representing a worst-case scenario for the search accuracy.

- In the other extreme case where $L = 1$, by Proposition 1, we must have $m = 1$ for $d(\mathbf{p}, \mathbf{q}) \leq 1$, $m \leq 1$ for $d(\mathbf{p}, \mathbf{q}) = 2$ and $m = 0$ for $d(\mathbf{p}, \mathbf{q}) > 2$. That is, for any string $\mathbf{p} \in \Omega$, the m value is either 0 or 1. We can determine for all strings with $m = 1$ that their Hamming distance to the query is not greater than two. For those with $m = 0$, they cannot be further ranked. Accordingly, if the Hamming distance of the true match is greater than two, which is most likely true in a biometric database, the inference-based search in this extreme case is also no better than a wild guess.

Intuitively, when L is neither too large nor too small, the values of the obfuscated distance measure m are more likely to span a wider range. As we have observed in Section 3.4 for the UBIRIS dataset, this makes it more feasible to order the retrieved candidates in multiple ranks and hence increases the search accuracy of the inference-based approach.

As a result of Proposition 1, for a pair of \mathbf{p} and \mathbf{q} with a particular value of m , the actual Hamming distance $d(\mathbf{p}, \mathbf{q})$ may take any integer value from a certain range. On the other hand, if $1 < L < D$, it is clear that different distributions of the mismatching bits between \mathbf{p} and \mathbf{q} can lead to different values of m even if the Hamming distance $d(\mathbf{p}, \mathbf{q})$ retains the same. Such a one-to-many relationship between the Hamming distance and the obfuscated distance measure ensures in all cases more or less privacy protection for the inference-based approach.

Fig. 7 illustrates the effect of distance obfuscation with $L = 30$ for the UBIRIS dataset. As discussed in Section 3.3, we consider a total of 3,840,000 pairwise comparisons. Among them, given the Hamming radius $r = 50$, we have 6,026 neighbouring pairs and 3,833,974 non-neighbouring pairs. In Fig. 7a, we plot the mapping from the Hamming distance d to the obfuscated distance measure m for all the 3,840,000 pairwise comparisons, which clearly exhibits a one-to-many relationship between d and m . In Fig. 7b, we present the histogram of d and that of m for all neighbouring pairs (i.e., those with $d \leq 50$). Likewise, the histograms for all non-neighbouring pairs (i.e., those with $d > 50$) are presented in Fig. 7c. In both cases, we observe that the distribution of d is very different from that of m .




Fig. 8. Retrieval performance of MIMP-RM on the simulated dataset with positions of mismatching bits randomly chosen from a range.

6 PERFORMANCE EVALUATION

In this section, we evaluate the performance of the proposed MIMP scheme in randomized Montgomery domains for inference-based similarity search. For ease of description, we shall hence call it as MIMP-RM for short. As discussed in Section 3.4, a simple implementation of k -NN search using MIMP-RM is to choose the top k candidates with the largest values of the obfuscated distance measure m . We demonstrate that the accuracy of candidate retrieval using this simple and privacy-preserving approach is close to that of conventional similarity search based on explicit distance values, but the associated cost is significantly reduced compared to cryptographic methods. We implement the test algorithms in Matlab and run the experiments on a 3.4 GHz Intel machine. For all experiments, the parameters used in MIMP-RM to generate the search indexes are $T = 2$, $c_N = 15$, and $c_R = 15$.

6.1 Impact of Mismatching Bit Distribution

Here, we investigate how the mismatching bit distribution can affect the retrieval performance of MIMP-RM. Specifically, we randomly generate a query set containing 200 binary strings of length $D = 2800$ bits. For each query \mathbf{q} , we randomly generate a database of 200 records in the following way. First, only ten of them are r -neighbours of the query \mathbf{q} where $r = 350$. Second, we consider three different distributions of the mismatching bits. That is, the positions of the mismatching bits are randomly chosen from the range $[1, D/i]$ for $i = 1, 2, 4$. In this way, a smaller i corresponds to more evenly distributed mismatching bits over the binary string. Given that we know the ground truth set $\mathcal{B}(\mathbf{q}; r)$ in this experiment, we perform MIMP-RM with $L = 200$ on the simulated dataset and evaluate the quality of the top k retrieved candidates.

Fig. 8 plots the *precision-recall* (PR) curves under the three different mismatching bit distributions. The PR curves are derived by varying k from 1 to 30. The results demonstrate that both the precision and the recall can be significantly improved if the mismatching bits are more evenly distributed. Thus, in cases where the mismatching bits are less evenly distributed, to improve the retrieval performance based on matching piecewise binary codes, it is helpful to apply a random synchronized permutation to binary feature vectors before performing the search. The results also show that the precision is very high when the recall is low, corresponding to the cases where k is small and indicating that the

retrieved candidates are mostly relevant. This further demonstrates the effectiveness of inference-based r -neighbour detection where the retrieved candidates with large m values have a high probability that they are r -neighbours of the query as discussed in Section 3.3.

6.2 Impact of the Parameter L

In Section 5, we argued that the number of substrings L is an important parameter of MIMP-RM. In particular, we showed in Section 5.2 that, to reduce the information leakage in randomized Montgomery domains, it is desirable to choose a large value of s , corresponding to a small value of L . In Section 5.4, we established that choosing an appropriate value for L is also subject to a fundamental trade-off between privacy and search accuracy. Here, we provide more extensive results and further demonstrate the impact of L on the search accuracy of MIMP-RM.

The experiments are conducted on public biometric datasets. In addition to the UBIRIS dataset, we also consider two benchmark face datasets. In particular, the LFW dataset [43] includes 13,233 face images crawled from the Internet for 5,749 subjects, and the FERET dataset [44] contains 2,400 face images of 200 subjects each with 12 images taken under semi-controlled environment.

Nowadays, face recognition based on features learned from deep convolution networks, known as deep features, can achieve an accuracy over 99 percent on the once hardest LFW dataset [3]. For the LFW dataset, we leverage the state-of-the-art deep features learned from a CNN framework [45]. The deep network is trained on the CASIA-WebFace dataset [46]. Here, we use LSH embedding [47] to convert the deep features into a binary string representation of 400 bits per LFW face image. To test our method with different data characteristics, we employ supervised PCA features on the FERET dataset. Specifically, we use six face samples per identity for training and six face samples for testing. We apply linearly separable subcodes [14] to convert the less robust PCA features into binary string representations of 448 bits per FERET face image, followed by a random synchronized permutation.

We enrol one sample per identity in the database and use another sample of that identity to search in each run of the experiments. The search accuracy is evaluated in terms of the *hit rate*, defined as the percentage of queries with true match found in the top k retrieved candidates. Fig. 9 presents the search accuracy results obtained by varying k from 1 to 30. We test four different values of L , i.e., 25, 50, 75 and 100, respectively.

For all the three datasets, we observe that the hit rate increases significantly from $L = 25$ to $L = 50$ and drops at $L = 100$, which is consistent with discussions in Section 5.4. We also observe that the hit rate seems to peak at $L = 75$ since there is only a marginal improvement from $L = 50$ to $L = 75$. To understand this latter effect, we perform an empirical study on the test samples of each database, and it reveals that over 99 percent of true match pairs have a Hamming distance smaller than $r^* = 150$. Recall from Proposition 2 that, in this case, with $L > r^*/2 = 75$, all such true match pairs are guaranteed to have a collision count and thus can be detected. Accordingly, in the context of MIMP-RM, choosing $L = r^*/2$ may serve as a rule of thumb for




Fig. 9. Impact of the parameter L on the search accuracy of MIMP-RM with respect to the top k retrieved candidates. (a) The UBIRIS dataset. (b) The LFW dataset. (c) The FERET dataset.

maximizing the search accuracy, and one may use a smaller L value to strike a balance between search accuracy and privacy while reducing the information leakage of hash-based indexing in randomized Montgomery domains.

6.3 Search Accuracy

We demonstrate the effectiveness of MIMP-RM in search accuracy by comparing it with two baseline approaches that can be used in the context of privacy-preserving similarity search. In addition to the “LSH + partial distance” approach [29], we also consider what we call the “exhaustive Hamming” approach that performs exhaustive search by pairwise comparison based on full-string representations in Hamming space. Note that “exhaustive Hamming” provides the highest search accuracy that is achievable with secure computation methods as the latter can retain one-to-one matching accuracy in the encrypted domain [22]. However, such cryptography-based methods require pairwise comparison in the encrypted domain and have high computation cost as will be shown in Section 6.4.

For “LSH + partial distance”, we generate the unencrypted part by random sampling of bit positions as in Section 5.3. In this way, we create L hash tables each being indexed by sub-hash codes of s bits. To increase the ambiguity of query information for the server [28], each sub-hash code of the query has one bit value omitted so that the server needs to fill the absent bits. The retrieved candidates from the hit buckets are then ranked in the ascending order of their partial distance computed from the plaintext bit positions, i.e., the unencrypted part of the full-string representation.

Fig. 10 plots the search accuracy of the comparing methods in terms of the hit rate with respect to the k value. The parameter L is set to 50, 50 and 56 for the UBIRIS dataset, the LFW dataset and the FERET dataset, respectively, which results in $s = 8$ in all cases. We observe that MIMP-RM

performs closely to “exhaustive Hamming”, especially in Figs. 10a and 10b where the full-string representations are highly discriminative in Hamming space. In all cases, MIMP-RM outperforms “LSH + partial distance” significantly. We believe it is mainly because the LSH-based approach in general has a low recall especially when the number of random hash functions is relatively small. In the above setting, the average number of plaintext bit positions is 233, 233 and 261 for the three datasets, respectively, representing about 58 percent of the full-string length. In contrast, MIMP-RM utilizes the full-string binary feature but in an obfuscated way that can provide highly accurate search results while increasing the uncertainty for adversarial learning in randomized Montgomery domains.

6.4 Computation Cost in Montgomery Domains

Index generation and search in our proposed MIMP-RM approach mainly involve computing nested Montgomery forms. As discussed in Section 4.1, the Montgomery form of a natural number can be evaluated efficiently via (15). To demonstrate the efficiency of computation in Montgomery domains, we compare it with two popular cryptographic methods, namely Paillier homomorphic encryption and the cryptographic hash function of SHA-256. We implement the former with the homomorpheR library and the latter with the sodium library on an R platform.

Table 2 provides the computation cost of each method conducted in R/3.2.2. The results are obtained for encoding piecewise binary codes that are generated from the UBIRIS dataset. The resulting cipher codes are further compressed using ZIP. We evaluate the computation cost by four measures. In particular, “Ciphertext size” and “Zipped file size” refer to the file size before and after compression, respectively. “Encoding time” refers to the CPU time required for encoding all records in the database. “Scanning time” refers




Fig. 10. Performance comparison in terms of the search accuracy with respect to the top k retrieved candidates. (a) The UBIRIS dataset. (b) The LFW dataset. (c) The FERET dataset.

TABLE 2
Comparison of Computation Cost for Encoding Piecewise Binary Codes Generated from the UBIRIS Dataset

	Ciphertext size (MB)	Zipped file size (MB)	Encoding time (sec)	Scanning time (sec)
Paillier homomorphic encryption	24	3.8	168.98	301.03
Cryptographic hash function (SHA-256)	3	0.437	81.72	140.26
Nested Montgomery form	0.155	0.029	1.01	0.0045

to the CPU time required for a linear scan of the database for pairwise matching.

It can be seen in Table 2 that randomized signatures generated in Montgomery domains have much smaller code size and higher processing speed. Using the nested Montgomery forms, encoding is over 150 times faster and scanning is up to five orders of magnitude faster than that using Paillier homomorphic encryption. The results confirm that the efficiency of computation is significantly improved when done in Montgomery domains.

7 CONCLUSION

We have studied the problem of privacy-preserving similarity search in a biometric database to withstand adversarial machine learning based on the critical biometric similarity information. Unlike existing privacy protection methods for biometric identification that are in general required to perform cumbersome distance comparisons in the encrypted domain, the new approach proposed in this paper is based on statistical inference and carefully designed data obfuscation mechanisms that obviate the need for comparing exact distance values. We have also proposed to protect the data structures by performing hash-based indexing in randomized Montgomery domains with virtually negligible information leakage. Experiments on public biometric datasets have confirmed that candidate retrieval using our simple and privacy-preserving approach is both statistically reliable and computationally efficient. In future work, we are interested in extending the inference-based approach to other biometric modalities such as fingerprints that do not have a fixed-length binary string representation. We are also interested in extending the inference-based approach to search problems with distance measures other than in Hamming space.


ACKNOWLEDGMENTS

The authors wish to thank the anonymous reviewers for their valuable comments that contributed to the improved quality of this paper. This work is supported by the Research Grants Council of Hong Kong (HKBU12202214) and the National Natural Science Foundation of China (61403324 and 61502058). Most of the work was done while Y. Wang was with Hong Kong Baptist University.

REFERENCES


- [1] Iris recognition, palm-vein, fingerprinting: Which biometric is best for healthcare?, *Healthcare IT News*, Mar. 2016. [Online]. Available: <http://www.healthcareitnews.com/news/iris-scanning-palm-vein-fingerprinting-which-biometric-best-healthcare>
- [2] Indian bank introduces Aadhaar-based ATM, Apr. 2016. [Online]. Available: <http://www.biometricupdate.com/tag/atm>
- [3] D. Wang, C. Otto, and A. K. Jain, "Face search at scale," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1122–1136, Jun. 2017.
- [4] K. Nandakumar and A. K. Jain, "Biometric template protection: Bridging the performance gap between theory and practice," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 88–100, Sep. 2015.
- [5] J. Bringer, H. Chabanne, and A. Patey, "Identification over out-sourced biometric data," in *Biometric Security*. Newcastle upon Tyne, U.K.: Cambridge Scholars Publishing, 2015, ch. 10, pp. 312–350.
- [6] B. Biggio, G. Fumera, P. Russu, L. Didaci, and F. Roli, "Adversarial biometric recognition: A review on biometric system security from the adversarial machine-learning perspective," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 31–41, Sep. 2015.
- [7] J. Galbally, C. McCool, J. Fierrez, S. Marcel, and J. Ortega-Garcia, "On the vulnerability of face verification systems to hill-climbing attacks," *Pattern Recognit.*, vol. 43, no. 3, pp. 1027–1038, Mar. 2010.
- [8] T. Boulton, W. Scheirer, and R. Woodworth, "Revocable fingerprint biotokens: Accuracy and security analysis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2007, pp. 1–8.
- [9] R. N. Rodrigues, L. L. Ling, and V. Govindaraju, "Robustness of multimodal biometric fusion methods against spoof attacks," *J. Vis. Languages Comput.*, vol. 20, no. 3, pp. 169–179, Jun. 2009.
- [10] B. Biggio, G. Fumera, G. L. Marcialis, and F. Roli, "Statistical meta-analysis of presentation attacks for secure multibiometric systems," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 3, pp. 561–575, Mar. 2017.
- [11] Y. Wang, L. Wang, Y.-M. Cheung, and P. C. Yuen, "Learning binary compact codes for hash-based fingerprint indexing," *IEEE Trans. Inf. Forensics Secur.*, vol. 10, no. 8, pp. 1603–1616, Aug. 2015.
- [12] J. Lu, V. E. Liang, X. Zhou, and J. Zhou, "Learning compact binary face descriptor for face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 10, pp. 2041–2056, Oct. 2015.
- [13] Z. Sun and T. Tan, "Ordinal measures for iris recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 12, pp. 2211–2226, Dec. 2009.
- [14] M.-H. Lim and A. B. J. Teoh, "A novel encoding scheme for effective biometric discretization: Linearly separable subcode," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 2, pp. 300–313, Feb. 2013.
- [15] M. Norouzi, A. Punjani, and D. Fleet, "Fast exact search in Hamming space with multi-index hashing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 6, pp. 1107–1119, Jun. 2014.
- [16] M. Muja and D. G. Lowe, "Scalable nearest neighbor algorithms for high dimensional data," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 11, pp. 2227–2240, Nov. 2014.
- [17] Y. Wang and J. Hu, "Global ridge orientation modeling for partial fingerprint identification," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 1, pp. 72–87, Jan. 2011.
- [18] Y. Wang, J. Hu, and D. Phillips, "A fingerprint orientation model based on 2D Fourier expansion (FOMFE) and its application to singular-point detection and fingerprint indexing," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 4, pp. 573–585, Apr. 2007.
- [19] J. Bringer, H. Chabanne, and A. Patey, "Privacy-preserving biometric identification using secure multiparty computation: An overview and recent trends," *IEEE Signal Process. Mag.*, vol. 30, no. 2, pp. 42–52, Mar. 2013.
- [20] M. Barni, G. Droandi, and R. Lazzeretti, "Privacy protection in biometric-based recognition systems: A marriage between cryptography and signal processing," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 66–76, Sep. 2015.
- [21] M. Haghghat, S. Zonouz, and M. Abdel-Mottaleb, "CloudID: Trustworthy cloud-based and cross-enterprise biometric identification," *Expert Syst. Appl.*, vol. 42, no. 21, pp. 7905–7916, Jun. 2015.
- [22] W. Lu, A. L. Varna, and M. Wu, "Confidentiality-preserving image search: A comparative study between homomorphic encryption and distance-preserving randomization," *IEEE Access*, vol. 2, pp. 125–141, Mar. 2014.
- [23] S. Rane, Y. Wang, S. Draper, and P. Ishwar, "Secure biometrics: Concepts, authentication architectures, and challenges," *IEEE Signal Process. Mag.*, vol. 30, no. 5, pp. 51–64, Sep. 2013.
- [24] V. Patel, N. Ratha, and R. Chellappa, "Cancelable biometrics: A review," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 54–65, Sep. 2015.
- [25] A. Cavoukian and A. Stoianov, "Privacy by design solutions for biometric one-to-many identification systems," *Information and Privacy Commissioner*, Toronto, ON, Canada, Jun. 2014, <https://www.ipc.on.ca/wp-content/uploads/2014/06/pbd-solutions-biometric.pdf>

- [26] R. Lagendijk, Z. Erkin, and M. Barni, "Encrypted signal processing for privacy protection: Conveying the utility of homomorphic encryption and multiparty computation," *IEEE Signal Process. Mag.*, vol. 30, no. 1, pp. 82–105, Jan. 2013.
- [27] S. Rane and P. Boufounos, "Privacy-preserving nearest neighbor methods: Comparing signals without revealing them," *IEEE Signal Process. Mag.*, vol. 30, no. 2, pp. 18–28, Mar. 2013.
- [28] L. Weng, L. Amsaleg, A. Morton, and S. Marchand-Maillet, "A privacy-preserving framework for large-scale content-based information retrieval," *IEEE Trans. Inf. Forensics Secur.*, vol. 10, no. 1, pp. 152–167, Jan. 2015.
- [29] L. Weng, L. Amsaleg, and T. Furon, "Privacy-preserving outsourced media search," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 10, pp. 2738–2751, Oct. 2016.
- [30] M. Gertz and S. Jajodia, *Handbook of Database Security - Applications and Trends*. Berlin, Germany: Springer, 2008.
- [31] D. Bakken, R. Rameswaran, D. Blough, A. Franz, and T. Palmer, "Data obfuscation: Anonymity and desensitization of usable data sets," *IEEE Secur. Privacy*, vol. 2, no. 6, pp. 34–41, Nov. 2004.
- [32] A. Jiménez, B. Raj, J. Portelo, and I. Trancoso, "Secure modular hashing," in *Proc. IEEE Workshop Inf. Forensics Secur.*, 2015, pp. 1–6.
- [33] A. Aghasaryan, M. Bouzid, D. Kostadinov, and M. Kothari, "On the use of LSH for privacy preserving personalization," in *Proc. IEEE Int. Conf. Trust Secur. Privacy Comput. Commun.*, 2013, pp. 362–371.
- [34] P. Boufounos and S. Rane, "Secure binary embeddings for privacy preserving nearest neighbors," in *Proc. IEEE Workshop Inf. Forensics Secur.*, Nov. 2011, pp. 1–6.
- [35] A. Aghasaryan, M. Bouzid, D. Kostadinov, and A. Nandi, "Exploring the impact of LSH parameters in privacy-preserving personalization," *Bell Labs Tech. J.*, vol. 18, no. 4, pp. 33–44, Mar. 2014.
- [36] H. V. Poor, *An Introduction to Signal Detection and Estimation*. Berlin, Germany: Springer-Verlag, 1994.
- [37] J. A. Rice, *Mathematical Statistics and Data Analysis*, 3rd ed. Belmont, CA, USA: Cengage Learning, 2007.
- [38] University of Bath Iris Image Database, 2009. [Online]. Available: <http://www.smartsensors.co.uk/irisweb/database.htm>
- [39] R. Crandall and C. Pomerance, *Prime Numbers: A Computational Perspective*. New York, NY, USA: Springer, 2001.
- [40] J. Hromkovic, *Design and Analysis of Randomized Algorithms: Introduction to Design Paradigms*. Berlin, Germany: Springer, 2005.
- [41] L. Sweeney, "K-anonymity: A model for protecting privacy," *Int. J. Uncertainty Fuzziness Knowl.-Based Syst.*, vol. 10, pp. 557–570, 2002.
- [42] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed. New York, NY, USA: Wiley-Interscience, 2006.
- [43] E. Learned-Miller, G. B. Huang, A. RoyChowdhury, H. Li, and G. Hua, "Labeled faces in the wild: A survey," in *Advances in Face Detection and Facial Image Analysis*. Berlin, Germany: Springer, 2016, pp. 189–248.
- [44] P. Phillips, H. Moon, S. Rizvi, and P. Rauss, "The FERET evaluation methodology for face-recognition algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 10, pp. 1090–1104, Oct. 2000.
- [45] X. Wu, "Learning robust deep face representation," *arXiv:1507.04844*, 2015, <https://arxiv.org/pdf/1507.04844.pdf>
- [46] D. Yi, Z. Lei, S. Liao, and S. Z. Li, "Learning face representation from scratch," *arXiv:1411.7923*, 2014, <https://arxiv.org/pdf/1411.7923.pdf>
- [47] A. Andoni and P. Indyk, "Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions," *Commun. ACM*, vol. 51, no. 1, pp. 117–122, Jan. 2008.




Yi Wang received the PhD degree in computer science from RMIT University, Melbourne, Australia, in 2009. She was a research associate in the School of Mathematics and Statistics, The University of New South Wales, Sydney, in 2009–2012, and a research assistant professor in the Department of Computer Science, Hong Kong Baptist University, in 2012–2016. In 2017, she entered the China's "1000-talents Plan" program for young researchers and joined Dongguan University of Technology, China, as an associate

professor. Her research interests include biometrics, pattern recognition, and privacy-aware computing.



Jianwu Wan received BS, MS, and PhD degrees from the School of Computer Science and Technology, Nanjing Normal University, in 2008, 2010, and 2013, respectively. Then, he joined the School of Information Science and Engineering, Changzhou University. In 2015–2016, he was a postdoctoral research fellow in the Department of Computer Science, Hong Kong Baptist University. His research interests include machine learning and pattern recognition.



Jun Guo received the PhD degree in electrical and electronic engineering from The University of Melbourne, Australia, in 2006. He was in the School of Computer Science and Engineering, The University of New South Wales, Australia, as a senior research associate from 2006 to 2008 and on an Australian postdoctoral fellowship supported by the Australian Research Council from 2009 to 2011. From 2012 to 2016, he was in the Department of Electronic Engineering, City University of Hong Kong. He is currently an associate professor in the School of Computer Science and Network Security, Dongguan University of Technology, China. His current research interests include networking and security.



Yiu-Ming Cheung received the PhD degree from the Department of Computer Science and Engineering, Chinese University of Hong Kong. He is a full professor in the Department of Computer Science, Hong Kong Baptist University. His current research interests focus on machine learning, pattern recognition, visual computing, and optimization. He is an IET/IEE fellow and BCS fellow. More details can be found at: <http://www.comp.hkbu.edu.hk/~ymc>.



Pong C. Yuen received the PhD degree in electrical and electronic engineering from The University of Hong Kong, in 1993. He joined the Hong Kong Baptist University in 1993 and, currently is a professor and head of the Department of Computer Science. He has been actively involved in many international conferences and professional community. He was the track co-chair of the International Conference on Pattern Recognition (ICPR) 2006, the program co-chair of the IEEE Fifth International Conference on Biometrics:

Theory, Applications and Systems (BTAS) 2012 as well as the IEEE International Conference on Identity, Security and Behavior Analysis (ISBA) 2016. He serves as the program co-chair of International Conference on Pattern Recognition and Artificial Intelligence in 2018. Currently, he is the vice president (Technical Activities) of the IEEE Biometrics Council, editorial board member of the *Pattern Recognition*, associate editor of the *IEEE Transactions on Information Forensics and Security*, and senior editor of the *SPIE Journal of Electronic Imaging*. He also serves as a Hong Kong Research Grant Council Engineering Panel Member. His current research interests include video surveillance, human face recognition, biometric security and privacy.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.