



Non-interactive Privacy-Preserving Naïve Bayes Classifier Using Homomorphic Encryption

Jingwei Chen^{1,2}, Yong Feng^{1,2}, Yang Liu^{3(✉)}, Wenyuan Wu^{1,2},
and Guanci Yang⁴

¹ Chongqing Key Laboratory of Automated Reasoning and Cognition, Chongqing Institute of Green and Intelligent Technology, Chinese Academy of Sciences, Chongqing, China

{chenjingwei,yongfeng,wuwenyuan}@cigit.ac.cn

² Chongqing College, University of Chinese Academy of Sciences, Chongqing, China

³ Information Science and Engineering, Chongqing Jiaotong University, Chongqing, China

liuyang13@cqjtu.edu.cn

⁴ Key Laboratory of Advanced Manufacturing Technology of Ministry of Education, Guizhou University, Guiyang, China

guanci_yang@163.com

Abstract. In this paper, we propose a privacy-preserving naive Bayes classifier based on a leveled homomorphic encryption scheme due to Brakerski-Gentry-Vaikuntanathan (BGV). The classifier runs on a server that is also the owner of the model, with input as BGV encrypted data from a client. The classifier produces encrypted classification results which can only be decrypted by the client, whereas the model is only accessible to the server itself. This ensures that the classifier does not leak any private information on either the model of the server or the data and results of the client. More importantly, the classifier does not require any interaction between the server and the client during the classification phase. The main technical ingredient is an algorithm to compute the index of the maximum of an encrypted array homomorphically, which does not require any interaction. The proposed classifier is implemented using a homomorphic encryption library HELib. Preliminary experiments demonstrate the efficiency and accuracy of the proposed privacy-preserving naive Bayes classifier.

Keywords: Privacy-preserving data mining · Homomorphic encryption · Naïve Bayes classifier

This research was supported in part by National Key R&D Program of China (2020YFA0712303), NSFC (61903053), Guizhou Science and Technology Program [2020]4Y056, Chongqing Science and Technology Program (cstc2021jcyj-msxmX0821, cstc2018jcyj-yszxX0002, cstc2019yszx-jcyjX0003, cstc2020yszx-jcyjX0005, cstc2021yszx-jcyjX0004, KJQN201900702), and Youth Innovation Promotion Association of CAS (2018419).

© ICST Institute for Computer Sciences, Social Informatics and Telecommunications Engineering 2022

Published by Springer Nature Switzerland AG 2022. All Rights Reserved

W. Shi et al. (Eds.): SPNCE 2021, LNICST 423, pp. 192–203, 2022.

https://doi.org/10.1007/978-3-030-96791-8_14

1 Introduction

Over the past decade, Machine Learning as a Service (MLaaS) has been involved in various fields, from academia to industry. A typical application scenario is that the model vendor uses a large amount of user data to train the model and then uses the trained model to infer/predict some results based on data supplied by clients. However, as security incidents such as data breaches continue to occur, the demand for privacy-preserving MLaaS is rapidly increasing. On the one hand, the model owner is unwilling to leak information about the model. On the other hand, the data owner is reluctant to leak information about the data as well. To resolve this contradiction, privacy-preserving machine learning is proposed.

In this paper, we consider the framework presented in [2] for privacy-preserving classifiers. As shown in Fig. 1, each shaded box indicates private data that should be accessible to only one party: the model to the server, and the data and prediction result to the client. In particular, we present a privacy-preserving naïve Bayes classifier (Protocol 1) based on a leveled homomorphic encryption (LHE) scheme BGV [4], which is based on R-LWE [20] and hence thought to be post-quantum safe, and which allows us to evaluate functions with a bounded multiplicative depth on encrypted data.




Fig. 1. Framework of privacy-preserving classifiers

In Protocol 1, the data owner (client) encrypts its data x to be predicted as a ciphertext c and sends c to the model owner (server). After receiving the ciphertext c , the server evaluates the model using the client's public key, with the input as the encrypted data c . The server sends the resulting ciphertext to the client, and the client decrypts the ciphertext using itself's secret key to specify in which class x lies. Thanks to an algorithm to compute a ciphertext of the index of the maximum of an encrypted array (Algorithm 3), no interaction between the client and the server happens during the classification phase (Step 3–6 in Protocol 1). This is one of the most important feature of our classifier.

In Sect. 2, we give a brief introduction to homomorphic encryption and present several building blocks for our classifier, including the main technical ingredient, Algorithm 3. In Sect. 3, we propose a privacy-preserving naïve Bayes classifier and prove its correctness and its security in the passive

(or honest-but-curious [13]) model. In Sect. 4, we report some preliminary experimental results based on our implementation of Protocol 1.

Related Work. It seems to be impossible to list all literature on privacy-preserving protocols for classifiers. We refer to [2, Sec. II] for a good survey. Here we focus mainly on those privacy-preserving naïve Bayes classifiers based on homomorphic encryption (HE).

Naïve Bayes classifiers is a simple but powerful algorithm to predict the category label of unclassified samples; see, e.g., [7]. Bost *et al.* proposed in [2, Sec. VI] the first efficient privacy-preserving protocols for naïve Bayes classifier based on the Paillier encryption scheme [21], which would be broken by quantum attackers. Li *et al.* proposed in [18] a secure naïve Bayes classifier for four parties. Later on, Kim *et al.* [17] adapted Li *et al.*'s framework using the HE scheme BGV. Yasumura *et al.* [25] and Sun *et al.* [23] also give privacy-preserving protocols for naïve Bayes classification based on BGV. All of these privacy-preserving naïve Bayes classifiers require interaction between participants. During the classification phase, Protocol 1 presented in this paper does not require interaction at all. Furthermore, being different from those protocols based on the Paillier scheme, our protocol is post-quantum safe because of the post-quantum safe BGV.

2 Preliminaries

In this section, we give some basics on homomorphic encryption schemes, which will be helpful for the rest of the paper. We refer the reader to [4, 11, 14] for more details. We also introduce a few building blocks for our protocol.

2.1 Homomorphic Encryption

A *homomorphic encryption* (HE) scheme is composed of four algorithms: KeyGen, Enc, Dec and Eval. Given a security parameter λ as input, KeyGen outputs a secret key sk , a public key pk , and an evaluation key ek . The encryption algorithm encrypts a plaintext b to a ciphertext $c \leftarrow \text{Enc}_{pk}(b)$, and the decryption algorithm decrypts a ciphertext c to a plaintext $b := \text{Dec}_{sk}(c)$. The evaluation algorithm evaluates a function of f on a ciphertext $c = \text{Enc}_{pk}(b)$ to a new ciphertext $c' \leftarrow \text{Eval}_{ek}(f, c)$ satisfying $\text{Dec}_{sk}(c') = f(b)$. For convenience, we call f an *HE function*.

For almost all known HE schemes, every ciphertext has noise, and every operation on ciphertext introduces noise in the resulting ciphertext. A *Leveled HE* (LHE) scheme allows to evaluate any HE function f with bounded number of arithmetic/boolean operations. Beyond this bound, such an LHE scheme cannot correctly decrypt the ciphertext since the accumulated noise is too large. Contrarily, if an HE scheme supports to evaluate an arbitrary computable function, then it is called a *fully homomorphic encryption* (FHE) scheme, which depends on the so-called *bootstrapping* technique [12] to reduce noise.

Since Gentry’s seminal work [12], multiple HE schemes have been designed, such as BGV [4], BFV [3, 10], CKKS [5], FHEW [9], TFHE [6]. Each of them has its features. For instance, BGV and BFV are good at performing large vectorial arithmetic operations, CKKS supports floating-point computations and FHEW and TFHE run bootstrapping for one bit extremely fast but slow for arithmetic operations. Since naïve Bayes classifiers require many integer arithmetic operations, and as reported in [1, 19] that BGV is faster than BFV, especially for cases with large modulus, we choose an LHE variant of the BGV scheme, which will be introduced in the next subsection.

2.2 The BGV Scheme

For completeness, we give a brief description of a variant of BGV that is implemented in HELib [15].

- **Setup**(1^λ). Given a security parameter λ as input, set an integer m (that defines the m -th cyclotomic polynomial $\Phi_m(x)$), a prime number p (the plaintext space is $R_p = \mathbb{Z}_p[x]/\langle\Phi_m(x)\rangle$), an odd modulus q (the ciphertext space is $R_q = \mathbb{Z}_q[x]/\langle\Phi_m(x)\rangle$), and a noise distribution χ over R_q . Output $params = (m, p, q, \chi)$.
- **KeyGen**($params$). Sample $s \leftarrow \chi$. Let $\mathbf{s} = (1, s) \in R_q^2$. Set $\mathbf{sk} = \mathbf{s}$. Generate $a \leftarrow R_q$ uniformly at random and a column vector with “small” coefficients $e \leftarrow \chi$. Set $b = a \cdot s + p \cdot e$. Output $\mathbf{sk} = \mathbf{s}$ and the public key $\mathbf{a} = (b, -a)$.
- **Enc**($params, \mathbf{pk}, b$). To encrypt a message $b \in R_p$, set $\mathbf{m} = (b, 0) \in R_p^2$, sample a polynomial $r \leftarrow R$ with small coefficients and output the ciphertext $\mathbf{c} = \mathbf{m} + r\mathbf{a} \in R_q^2$.
- **Dec**($params, \mathbf{sk}, \mathbf{c}$). Output the message $b = [[\langle \mathbf{c}, \mathbf{s} \rangle]_q]_p$.

The quantity $\langle \mathbf{c}, \mathbf{s} \rangle$ is called the *noise* of the ciphertext \mathbf{c} under the secret key \mathbf{s} . Decryption works correctly only if the noise does not wrap around modulo q . If it is the case, the correctness follows from $[[\langle \mathbf{c}, \mathbf{s} \rangle]_q]_p = [[\langle \mathbf{m} + r\mathbf{a}, \mathbf{s} \rangle]_q]_p = [b + p \cdot re]_p = b$. The security is based on the Ring-LWE assumption [20], which is thought to resist attacks from quantum computers.

Homomorphic Evaluation. The BGV scheme supports homomorphic addition and multiplication. Let \mathbf{c}_1 and \mathbf{c}_2 be two ciphertexts of two plaintexts b_1 and b_2 under the same secret key \mathbf{s} . Suppose that the noise of \mathbf{c}_1 and \mathbf{c}_2 is bounded from above by B . The addition (BGV.Add) of the two ciphertexts is simply $\mathbf{c}_+ = \mathbf{c}_1 + \mathbf{c}_2$, which is a ciphertext of $b_1 + b_2$ under the secret key \mathbf{s} . The noise of \mathbf{c}_+ is at most $2B$. For multiplication (BGV.Mul), $\mathbf{c}_\times = \mathbf{c}_1 \otimes \mathbf{c}_2$ is a ciphertext of $b_1 \cdot b_2$ under a new secret key $\mathbf{s} \otimes \mathbf{s}$, where \otimes is the usual tensor product. The noise of \mathbf{c}_\times can only be bounded from above by B^2 . To keep the secret key with small size and to decrease the noise of evaluated ciphertext, a refresh procedure BGV.Refresh (consisting of key switching and modulus switching) follows every homomorphic addition and multiplication. Of course, one can call BGV.Refresh only if necessary for efficiency. Note that the public key \mathbf{pk}

of BGV also includes all keys for BGV.Refresh. Theoretically, the cost of each homomorphic addition or multiplication increases fast as L grows, where L is the circuit depth of the function f to be evaluated; see, e.g., [4] for more details. Besides, BGV also supports plaintext-ciphertext addition (BGV.AddConst) and plaintext-ciphertext multiplication (BGV.MulConst).

Batching. Recall the plaintext space $R_p = \mathbb{Z}_p[x]/\langle \Phi_m(X) \rangle$. Let d be the multiplicative order of p modulo m , and $\phi(m)$ be the Euler’s totient function. Then d divides $\phi(m)$ and $R_p \cong \mathbb{F}_{p^d}^\ell$ with $\ell = \phi(m)/d$. Therefore each plaintext can be seen as a packed message with ℓ slots. From this view, each homomorphic operation on a ciphertext is equivalent to the same operation on all slots independently and simultaneously. This batching technique was introduced by Smart and Vercauteren in [22] and significantly decreased the amortized cost of Ring-LWE based HE schemes. For batching, BGV supports data packing (BGV.Encode), data rotating (BGV.Rotate), and data shifting (BGV.Shift). Based on these operations, one can build some advanced functions. For instance, BGV.TotalSum converts a ciphertext that encrypts (z_1, \dots, z_t) into a ciphertext that encrypts (y, \dots, y) with $y = \sum_{i=1}^t z_i$.

2.3 Building Blocks

We now describe a few necessary building blocks that will be used to build our classifier. Note that all the following algorithms will be executed on the server and that the owner of pk (the public key) in these algorithms is the client, not the server since we follow the framework given in Fig. 1.

Plaintext Matrix-Encrypted Vector Multiplication. Matrix-vector multiplication is fairly common in practice. Here we focus on plaintext matrix-encrypted vector multiplication. Given a matrix $\mathbf{A} \in \mathbb{Z}^{s \times t}$ and an encryption of a vector $\mathbf{z} \in \mathbb{Z}^t$, our goal is to obtain an encryption of \mathbf{Az} . We give two methods based on different method to encrypt a vector.

Naïve Encoding. To encrypt a vector $\mathbf{z} = (z_i)_{i \leq t} \in \mathbb{Z}^t$, one can encrypt each entry z_i of \mathbf{z} as a ciphertext. The resulting encryption of \mathbf{z} is a vector $\mathbf{c} = (c_i)_{i \leq t} \in R_q^t$ of ciphertexts, whose i -th entry c_i is a ciphertext of z_i . Then we have Algorithm 1.

Input: $\mathbf{c}' = (c'_i)_{i \leq t} \in R_q^t$ (c'_i encrypts the i th entry of $\mathbf{z} = (z_i)_{i \leq t}$), and public key pk ;
 $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}^{s \times t}$.
Output: $(c_i)_{i \leq s}$ with $c_i = \text{Enc}_{\text{pk}}(\sum_{j=1}^t a_{i,j} z_j)$.
 1: For $i = 1, \dots, s$ do the following:
 2: $c_i \leftarrow \text{Enc}_{\text{pk}}(0)$;
 3: For $j = 1, \dots, t$, update $c_i := \text{Add}_{\text{pk}}(c_i, \text{MulConst}_{\text{pk}}(a_{i,j}, c'_j))$.
 4: **return** $(c_i)_{i \leq s}$.

Algorithm 1: Naïve plaintext matrix-encrypted vector multiplication

Packed Encoding. Instead of the above elements-wise method, we can pack the vector $\mathbf{z} \in \mathbb{Z}^t$ into t slots of one plaintext, and encrypt it to only one ciphertext $c \in R_q$, which leads to Algorithm 2.

Input: $c \in R_q$ that encrypts $u \in R_p$ with $u = \text{Encode}(\mathbf{z})$, and public key pk ; $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}^{s \times t}$.
Output: $(c_i)_{i \leq s}$ with $c_i = \text{Enc}_{\text{pk}}(\sum_{j=1}^t a_{i,j} z_j)$.
 1: For $i = 1, \dots, s$ do the following:
 2: Encodes the i -th row \mathbf{a}_i of \mathbf{A} as $v_i = \text{Encode}(\mathbf{a}_i)$;
 3: Computes $c_i = \text{TotalSum}_{\text{pk}}(\text{MulConst}_{\text{pk}}(v_i, c))$.
 4: **return** $(c_i)_{i \leq s}$.

Algorithm 2: Packed plaintext matrix-encrypted vector multiplication

Comparison among Encrypted Integers. Comparison is a commonly used function in many applications. In this work, we use a very recent comparator presented by Iliashenko and Zucca in [16], which supports comparison operations for BGV (BGV.Comparator). Essentially, this method homomorphically evaluates the Lagrange interpolated polynomial of the less-than function over $S = [0, (p-1)/2]$ defined as follows:

$$\text{LT}_S(x, y) = \begin{cases} 1, & \text{if } 0 \leq x < y \leq (p-1)/2, \\ 0, & \text{if } 0 \leq y \leq x \leq (p-1)/2. \end{cases}$$

Input: $\mathbf{c} = (c_0, \dots, c_{t-1}) \in R_q^t$ (c_i encrypts the i th entry of $\mathbf{z} = (z_0, \dots, z_{t-1})$) and public key pk .
Output: A ciphertext $c \in R_q$ that encrypts the index of the maximal value of \mathbf{z} .
 1: Set $c \leftarrow \text{Enc}_{\text{pk}}(0)$. For $i = 0, \dots, t-1$ do the following:
 2: Compute $c' := \prod_{k=0, k \neq i}^{t-1} \text{AddConst}_{\text{pk}}(1, \text{MulConst}_{\text{pk}}(-1, \text{Comparator}_{\text{pk}}(c_j, c_k)))$.
 3: Update $c := \text{Add}_{\text{pk}}(c, \text{MulConst}_{\text{pk}}(j, c'))$.
 4: **return** c .

Algorithm 3: Encrypted index of maximum of an encrypted array

Furthermore, one can apply the less-than function LT to compute the arg max function of an encrypted array, where the arg max function returns an encryption of the index of the maximum of an array. As a matter of fact, for a given array $\mathbf{z} = (z_0, \dots, z_{t-1})$, we have

$$\arg \max_i (z_i)_{0 \leq i \leq t-1} = \sum_{j=0}^{t-1} j \cdot \prod_{k=0, k \neq j}^{t-1} (1 - \text{LT}(z_j, z_k)),$$

which results in Algorithm 3. Note that one would better use some recursive methods in practice to compute the encrypted product in Step 2 of Algorithm 3 for saving multiplicative depth.

3 Privacy-Preserving Naïve Bayes Classification

In this section, we first introduce the plaintext Naïve Bayes Classifier, and then adapt the classifier to the privacy-preserving setup and prove its correctness and security.

3.1 Naïve Bayes Classifier

Naïve Bayes classifier is based on the assumption that all features are conditional independent. Consider a data set with s categories $1, \dots, s$ and n features X_1, \dots, X_n , where each feature X_k has at most t different values $1, 2, \dots, t$. Under the conditional independence assumption, the classification of a sample $\mathbf{x} = (x_1, \dots, x_n)$ is

$$s^* = \arg \max_{i=1, \dots, s} \Pr[Y = i] \prod_{k=1}^n \Pr[X_k = x_k | Y = i],$$

where $\Pr[Y = i]$ is the probability that each class i occurs, i.e., the *prior probability*, and $\Pr[X_k = x_k | Y = i]$ is the probability of the k th feature X_k to be $x_k \in \{1, 2, \dots, t\}$ when \mathbf{x} belongs to category i , i.e., the *likelihood*. As in [2], we only deal with the case that the domain of the feature values (the x_i 's) is discrete and finite, so the $\Pr[X_k = x_k | Y = i]$'s are probability masses.

3.2 Preparing the Model

If the domain of the feature values is continuous, we first find a bound B on the values and then discretize them by splitting $[-B, B]$ into several equal intervals. For example, if the domain of the k th feature X_k is continuous on $[-1, 1]$, then one can discretize X_k as $X_k = 0$ if $X_k \in [-1, 0)$ and $X_k = 1$ if $X_k \in [0, 1]$. This discretization technique enables our classifier to deal with continuous features as well, possibly at the cost of decreasing the prediction accuracy.

For convenience, we can further simplify the feature values x_1, \dots, x_n to $1, \dots, t$. Furthermore, for numerical stability, we work with the logarithm of the probability:

$$s^* = \arg \max_{i=1, \dots, s} \left\{ \log \Pr[Y = i] + \sum_{k=1}^n \log \Pr[X_k = x_k | Y = i] \right\}, \quad (1)$$

where $x_k \in \{1, \dots, t\}$. Another convenient simplification is to take the numbering of the s classes as contiguous integers from 1 to s . Then s^* is precisely the index of the maximum over the s values in (1).

Additionally, since the BGV encryption scheme works with integers, one needs to convert each logarithm of probability in (1) to an integer by multiplying it with a certain number $K > 0$ and rounding it to the closest integer. A similar *shifting* technique is already used and analyzed in, e.g., [2, 24].

In summary, for a data set with s categories and n features (each feature has at most t different values), the prior probability in the model will be converted into a vector $\mathbf{b} = (b_1, \dots, b_s) \in \mathbb{Z}^s$, where b_i is obtained by rounding $K \cdot \log(\Pr[Y = i])$ for an appropriate scaling integer K . The likelihoods will be converted into n matrices $\mathbf{A}_k \in \mathbb{Z}^{s \times t}$ for $k = 1, \dots, n$, where the (i, j) -entry of \mathbf{A}_k is derived by rounding $K \cdot \log \Pr[X_k = j | Y = i]$ with the same integer K .

3.3 Privacy-Preserving Naïve Bayes Classifier

To resolve the privacy concerns, the client should only obtain the classification result s^* without learning any information about the prior probability and likelihood, and the server should learn nothing about the client's data \mathbf{x} .

The client has data $\mathbf{x} = (x_1, \dots, x_n)$ with $x_k \in \{1, \dots, t\}$ and wants the server to predict which class \mathbf{x} is in by using a naïve Bayes classifier without leaking any information about \mathbf{x} . One choice of the client is to encrypt \mathbf{x} using himself's public key. However, since \mathbf{x} is encrypted, the server cannot decide which entry of \mathbf{A}_k should be chosen. For instance, the first feature of \mathbf{x} is x_1 , i.e., $X_1 = x_1$. To access the information about $\Pr[X_1 = x_1 | Y = i]$ in \mathbf{A}_1 , we need to select the (i, x_1) entry of \mathbf{A}_1 . However, as the first entry of \mathbf{x} , x_1 is only available in encrypted form on the server-side. To get around this obstacle, one can encode the sample \mathbf{x} as a 0-1 matrix

$$\mathbf{X} = (\mathbf{e}_{x_1}, \dots, \mathbf{e}_{x_n}) \in \{0, 1\}^{t \times n}, \quad (2)$$

where \mathbf{e}_j is the t -dimensional vector whose j th entry is 1 and all others are 0. Now, to select the x_k -th row of a matrix $\mathbf{A}_k \in \mathbb{Z}^{s \times t}$ is just to compute $\mathbf{A}_k \mathbf{e}_{x_k}$. If \mathbf{e}_{x_k} is in encrypted form, then this is a plaintext matrix-encrypted vector multiplication discussed in Sect. 2.3.

Now we are ready to present our privacy-preserving naïve Bayes classifier as Protocol 1, which does not require any interaction between the server and the client during the classification phase (Step 3–6).

We prove the security of our protocol using the secure two-party computation framework for passive adversaries. Roughly speaking, a passive adversary tries to learn as much private information from the other party; however, this adversary follows the prescribed protocol faithfully.

Proposition 1. *Protocol 1 is correct. It is secure in the honest-but-curious model, assuming that the used LHE scheme is secure.*

Input of the client: A sample $\mathbf{x} = (x_1, \dots, x_n)$ to be predicted, the secret and public key sk and pk .

Input of the server: The model consisting of the likelihood information $(\mathbf{A}_k)_{k \leq n}$ and the prior information $(b_i)_{i \leq s}$, and the client's public key pk .

- 1: The client encode \mathbf{x} to a matrix \mathbf{X} as in (2).
- 2: The client encrypts the column vectors e_{x_k} of \mathbf{X} for $k = 1, \dots, n$ and sends these ciphertexts to the server.
- 3: The server do the following:
 - 4: For $i = 1, \dots, s$, set $c_i \leftarrow \text{Enc}_{\text{pk}}(0)$ and update $c_i := \text{AddConst}_{\text{pk}}(c_i, b_i)$.
 - 5: For $k = 1, \dots, n$, calling Algorithm 1 or 2 with input as the ciphertexts of e_{x_k} , \mathbf{A}_k and pk outputs $(c'_i)_{i \leq s}$. Update $c_i := \text{Add}_{\text{pk}}(c_i, c'_i)$ for $i = 1, \dots, s$.
 - 6: Calling Algorithm 3 with input as $\mathbf{c} = (c_i)_{i \leq s}$ and pk returns c .
- 7: The server sends c to the client.
- 8: The client decrypt c to $y = \text{Dec}_{\text{sk}}(c)$.

Protocol 1: Privacy-preserving naïve Bayes classifier

Proof (Sketch). The correctness follows from that what the server does is to evaluate the following procedure homomorphically:

- 1: Set $\mathbf{y} := \mathbf{b}$, the information of the prior probability.
- 2: For $k = 1, \dots, n$, set $\mathbf{y} := \mathbf{y} + \mathbf{A}_k \cdot e_{x_k}$.
- 3: **return** y as the index of the maximum of $\mathbf{y} = (y_i)_{0 \leq i \leq s-1}$.

From the server-side, the only message received is an encryption of the client's data \mathbf{x} , and the only message sent to the client is a ciphertext of the prediction result. From the client-side, \mathbf{x} is encrypted, the classification result is also encrypted, and the model is not accessible. So the security of Protocol 1 follows from that of the used LHE scheme. \square

4 Experiments

We have implemented Protocol 1 in C++ using HELib (v2.1.0) [15]. In this section, we will report the prediction accuracy and calculation time of our implementation. All experiments run serially (using only one thread) on a laptop with a Ubuntu 20.04 OS as Windows Subsystem for Linux, Intel i7-10750H CPU, and 16 GB RAM.

Data Set. In this experiment, Wisconsin Breast Cancer (WBC) data set in UCI Machine Learning Repository [8] was used. The dataset has 683 effective samples, classified into two categories, i.e., $s = 2$. There are nine features for each sample, and each feature may take at most ten different values, i.e., $n = 9$ and $t = 10$. Among these 683 samples, 478 samples are used for training (70%), and the remaining 205 samples are used to test.

Parameter Setting. For WBC, the scaling factor K in Sect. 3.2 is set to 1. This leads that the entries of the rounded logarithm of likelihood \mathbf{A}_k for $k = 1, \dots, n$ are integers between -6 and 0 , and the entries of the rounded logarithm of the prior probability \mathbf{b} are bounded by 2. Hence the resulting integers to be compared must be at most $6n + 2 = 54$, which implies that $p = 113$ is enough for our purpose. In addition, m is fixed to 12883. In this setting, each plaintext in R_p has $\ell = 3960$ slots.

Accuracy. According to our experiments, the classification accuracy of our implementation of Protocol 1 based on HELib is about 97%. Note that this accuracy is almost the same as the plaintext (unencrypted) naïve Bayes classifier.

Timing. In Step 5 of Protocol 1, there are two choices (Algorithm 1 and Algorithm 2) for plaintext matrix-encrypted vector multiplication. We test them all and record their performance in Table 1. The row named “naïve” (“packed” resp.) is the performance of Protocol 1 based on Algorithm 1 (Algorithm 2 resp.). The columns with “Ave.” is the average execution time for each sample. Table 1 shows that the naïve variant outperforms the packed variant. For comparison, Kim *et al.* [17] reported that their proposed privacy-preserving naïve Bayes classifier takes 17h40m on a single core for the Iris data set ($s = 3$, $n = 4$, $t = 5$, 30 samples to be classified) of UCI Machine Learning Repository [8].

Table 1. Performance of two variants of Protocol 1

	$\log q$	λ	Compare (s)	Total (s)	Ave. Compare (s)	Ave. Total (s)
Naïve	382	101	387.90	985.31	1.89	4.75
Packed	476	76	369.95	1931.60	1.75	9.42

5 Conclusion

In this paper, we attempt to design privacy-preserving classifier protocols in the client-server setting. The server owns the model, which should not be accessible to any other, and the client also needs to preserve the privacy of the data to be predicted. As a result, we propose a privacy-preserving naïve Bayes classifier (Protocol 1) based on the LHE scheme BGV. We show that the classifier is correct and secure in the honest-but-curious model. The main feature of our classifier is that it does not require any interaction between the client and the server during the classification phase. We are considering how to use SIMD to accelerate the classifier and experiment with more data sets. In addition, how to optimize the performance further using, e.g., parallel computing, is also an interesting problem.

References

1. Aguilar Melchor, C., Kilijian, M.-O., Lefebvre, C., Ricosset, T.: A comparison of the homomorphic encryption libraries HELib, SEAL and FV-NFLlib. In: Lanet, J.-L., Toma, C. (eds.) SECITC 2018. LNCS, vol. 11359, pp. 425–442. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-12942-2_32
2. Bost, R., Popa, R.A., Tu, S., Goldwasser, S.: Machine learning classification over encrypted data. In: Proceedings of the 22nd Annual Network and Distributed System Security Symposium, San Diego USA, 8–11 February 2015. The Internet Society (2015). <https://doi.org/10.14722/ndss.2015.23241>
3. Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical GapSVP. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 868–886. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32009-5_50
4. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. *ACM Trans. Comput. Theory* **6**(3), 13:1–13:36 (2014). <https://doi.org/10.1145/2633600>
5. Cheon, J.H., Kim, A., Kim, M., Song, Y.: Homomorphic encryption for arithmetic of approximate numbers. In: Takagi, T., Peyrin, T. (eds.) ASIACRYPT 2017. LNCS, vol. 10624, pp. 409–437. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-70694-8_15
6. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: TFHE: fast fully homomorphic encryption over the torus. *J. Cryptol.* **33**(1), 34–91 (2019). <https://doi.org/10.1007/s00145-019-09319-x>
7. Domingos, P., Pazzani, M.: On the optimality of the simple Bayesian classifier under zero-one loss. *Mach. Learn.* **29**(2), 103–130 (1997). <https://doi.org/10.1023/A:1007413511361>
8. Dua, D., Graff, C.: UCI machine learning repository (2017). <http://archive.ics.uci.edu/ml>
9. Ducas, L., Micciancio, D.: FHEW: bootstrapping homomorphic encryption in less than a second. In: Oswald, E., Fischlin, M. (eds.) EUROCRYPT 2015. LNCS, vol. 9056, pp. 617–640. Springer, Heidelberg (2015). https://doi.org/10.1007/978-3-662-46800-5_24
10. Fan, J., Vercauteren, F.: Somewhat practical fully homomorphic encryption. *Cryptology ePrint Archive* (2012). <https://eprint.iacr.org/2012/144>
11. Gentry, C.: A fully homomorphic encryption scheme. Ph.D. thesis, Stanford University, Stanford (2009). <https://crypto.stanford.edu/craig/craig-thesis.pdf>
12. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Mitzenmacher, M. (ed.) Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing, Bethesda, USA, 31 May–2 June 2009, pp. 169–178. ACM, New York (2009). <https://doi.org/10.1145/1536414.1536440>
13. Goldreich, O.: Foundations of Cryptography - Basic Applications. Cambridge University Press, Cambridge (2004)
14. Halevi, S., Shoup, V.: Design and implementation of HELib: a homomorphic encryption library. *Cryptology ePrint Archive* (2020). <https://eprint.iacr.org/2020/1481>
15. HELib: An implementation of homomorphic encryption. <https://github.com/homenc/HELib>. Accessed Aug 2021
16. Iliashenko, I., Zucca, V.: Faster homomorphic comparison operations for BGV and BFV. *Proc. Priv. Enhancing Technol.* **2021**(3), 246–264 (2021). <https://doi.org/10.2478/popets-2021-0046>

17. Kim, S., Omori, M., Hayashi, T., Omori, T., Wang, L., Ozawa, S.: Privacy-preserving Naive Bayes classification using fully homomorphic encryption. In: Cheng, L., Leung, A.C.S., Ozawa, S. (eds.) ICONIP 2018. LNCS, vol. 11304, pp. 349–358. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-04212-7_30
18. Li, X., Zhu, Y., Wang, J.: Secure Naïve Bayesian classification over encrypted data in cloud. In: Chen, L., Han, J. (eds.) ProvSec 2016. LNCS, vol. 10005, pp. 130–150. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-47422-9_8
19. Lou, Q., Feng, B., Fox, G.C., Jiang, L.: Glyph: fast and accurately training deep neural networks on encrypted data. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., Lin, H. (eds.) Proceedings of NeurIPS 2020, 34th Conference on Neural Information Processing Systems, Vancouver, Canada, NeurIPS (2020). <https://proceedings.neurips.cc/paper/2020/hash/685ac8cadc1be5ac98da9556bc1c8d9e-Abstract.html>
20. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. *J. ACM* **60**(6), 43:1–35 (2013). <https://doi.org/10.1145/2535925>
21. Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48910-X_16
22. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. *Des. Codes Crypt.* **71**(1), 57–81 (2012). <https://doi.org/10.1007/s10623-012-9720-4>
23. Sun, X., Zhang, P., Liu, J.K., Yu, J., Xie, W.: Private machine learning classification based on fully homomorphic encryption. *IEEE Trans. Emerg. Top. Comput.* **8**(2), 352–364 (2020). <https://doi.org/10.1109/TETC.2018.2794611>
24. Tschiatschek, S., Reinprecht, P., Mücke, M., Pernkopf, F.: Bayesian network classifiers with reduced precision parameters. In: Flach, P.A., De Bie, T., Cristianini, N. (eds.) ECML PKDD 2012. LNCS (LNAI), vol. 7523, pp. 74–89. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33460-3_10
25. Yasumura, Y., Ishimaki, Y., Yamana, H.: Secure naïve Bayes classification protocol over encrypted data using fully homomorphic encryption. In: Indrawan-Santiago, M., Pardede, E., Salvadori, I.L., Steinbauer, M., Khalil, I., Anderst-Kotsis, G. (eds.) Proceedings of the 21st International Conference on Information Integration and Web-Based Applications & Services, Munich, Germany, 2–4 December 2019, pp. 45–54. ACM, New York (2019). <https://doi.org/10.1145/3366030.3366056>