



Private naive bayes classification of personal biomedical data: Application in cancer data analysis



Alexander Wood^{a,b,*}, Vladimir Shpilrain^{d,e}, Kayvan Najarian^{b,c}, Delaram Kahrobaei^{a,f}

^a Department of Computer Science, The Graduate Center, CUNY, New York, NY, 10016, USA

^b The Department of Computational Medicine and Bioinformatics, University of Michigan Center for Integrative Research in Critical Care (MCIRCC), Ann Arbor, MI, 48109, USA

^c Emergency Medicine Department, University of Michigan, Ann Arbor, MI, 48109, USA

^d Department of Mathematics, The Graduate Center, CUNY, New York, NY, 10016, USA

^e Department of Mathematics, The City College of New York, New York, NY, 10031, USA

^f Department of Computer Science, Tandon School of Engineering, New York University, New York, NY, USA

ARTICLE INFO

Keywords:

cryptographic protocols

Data privacy

Fully homomorphic encryption

Medical information systems

Predictive models

ABSTRACT

Clinicians would benefit from access to predictive models for diagnosis, such as classification of tumors as malignant or benign, without compromising patients' privacy. In addition, the medical institutions and companies who own these medical information systems wish to keep their models private when in use by outside parties. Fully homomorphic encryption (FHE) enables computation over encrypted medical data while ensuring data privacy. In this paper we use private-key fully homomorphic encryption to design a cryptographic protocol for private Naive Bayes classification. This protocol allows a data owner to privately classify his or her information without direct access to the learned model. We apply this protocol to the task of privacy-preserving classification of breast cancer data as benign or malignant. Our results show that private-key fully homomorphic encryption is able to provide fast and accurate results for privacy-preserving medical classification.

1. Introduction

Cryptographic methods hide information while machine learning looks to discover information, yet the fields overlap in many applications [1]. Medical data in particular requires a high level of privacy and HIPAA constraints were designed in order to provide patients with a necessary level of confidentiality. On the other hand, these records represent a wealth of data that has already been collected by various research institutions and hospitals.

The ability to use this information without compromising the patients' privacy would streamline collaboration between researchers and clinicians. Privacy concerns make it difficult for researchers to access large amounts of medical data. However, training classification models on a small or single-source data set can lead to overfitting or over generalization of the model's accuracy during the testing stage. This results in a model which is not generalizable to new data [2].

The ability to perform private classification, where researchers test a trained classification model on an encrypted external database without revealing the model, would provide medical researchers with a method of verifying the generalizability of their models. Another

valuable application for private classification is for assisted decision making and diagnosis systems. The owners of these systems do not wish to share their models, and clinicians do not wish to share their patients' data points. However, access to a well-trained model could assist clinicians in diagnosing disease in their patients.

A well-known machine learning classifier is Naive Bayes, a model which classifies a new data point using probabilities computed during a training phase. Naive Bayes earns its "naive" title due to the strong independence assumptions it makes between the features of the data. This assumption yields a transparent and understandable method for classification which often outperforms more sophisticated methods [3]. Similarly, when diagnosing a patient, clinicians try to define conditionally independent attributes which could be indicative of a disease [4]. Naive Bayes has outperformed more sophisticated methods and state-of-the-art diagnosis systems on 5 out of 8 real-life medical data sets, including localization of a primary tumor, prediction of recurrence of breast cancer, and rheumatological diagnosis [5], as well as in applications predicting heart disease [4].

Fully-homomorphic encryption (FHE), which allows for computation of arbitrary functions over encrypted data, is one approach towards

* Corresponding author. Department of Computer Science, The Graduate Center, CUNY, New York, NY, 10016, USA.

E-mail address: anwood@med.umich.edu (A. Wood).

enabling private classification. Current FHE research focuses on public-key cryptosystems which involve public encryption key and a private decryption key. Some of these applications focus specifically on genomic computation, including edit distance [6], string matching [7,8], genomic tests such as ancestry and paternity [9], and other genomic tests [10,11]. Other research has focused on the task of private classification, including neural networks [12,13], decision trees [14], and Fisher's linear discriminant classifier [1].

In contrast, private key cryptosystems require prior knowledge of the encryption/decryption key(s). In other words, if multiple parties wish to perform decryption in a private-key setting, they must first exchange keys over a secure channel. While this is considered a disadvantage when the goal is purely communication, it is not necessarily a disadvantage within medical applications due to different privacy concerns [15]. In this paper we propose a fully homomorphic private key protocol for private Naive Bayes classification, a private argmax protocol, and test the protocol on publicly available breast cancer data [16]. We classify each data point with an average time of 0.40298 s.

1.1. Fully homomorphic encryption

Fully homomorphic encryption was first conceptualized as a 'privacy transformation' which would allow for computing functions on encrypted data without first having to decrypt the information [17]. This concept is known today as a fully homomorphic encryption scheme, formally defined as an encryption scheme which allows for computation of arbitrary functions over encrypted data. Because a boolean circuit can be used to describe arbitrary computations, a scheme only needs to be homomorphic over addition and over multiplication in order to be described as fully homomorphic [18]. A scheme is called additively homomorphic if

$$[[x + y]] = [[x]] \oplus [[y]]$$

and multiplicatively homomorphic if

$$[[x \cdot y]] = [[x]] \otimes [[y]]$$

for operations \oplus, \otimes in the ciphertext space, where $[[\alpha]]$ denotes the encryption of α . While computation of arbitrary functions is *theoretically* possible using addition and multiplication as described above, this in itself is not sufficient for *practical* computation of arbitrary functions. Any fully homomorphic encryption scheme which is suited for practical use will only be able to compute polynomial functions and polynomial approximations of functions. This has been termed *polynomial machine learning* in the context of fully homomorphic encryption [1].

Gentry's first fully homomorphic encryption scheme is lattice-based [19]. Improvements on this method led quickly to the second generation of fully homomorphic encryption schemes [20–22] including the BGV scheme [20,21,23] and YASHE [24]. Recent improvements have built upon this foundation to yield even faster schemes which allow for practical applications [25–30].

1.2. GKS FHE scheme

The Gribov-Kahrobaei-Shpilrain (GKS) scheme is a ring-based private-key FHE scheme is able to avoid much of the computational overhead required for fully homomorphic public key encryption and is secure against ciphertext-only attack [15]. Define the ring S_n as

$$S_n = \langle x_1, x_2, \dots, x_n \mid p \cdot 1 = 0, x_i^2 = x_i, \text{ and } x_i x_j = x_j x_i \text{ for all } i, j \rangle.$$

This ring contains a super-exponential (in n) number of *idempotent* elements, or elements $g \in S_n$ such that $g^2 = g$. Some idempotent elements are given by

$$e_F = \prod_{i \in F} x_i \cdot \prod_{j \notin F} (1 - x_j)$$

for any set of indexes $F \subseteq \{1, 2, \dots, n\}$. The above construction yields 2^n idempotents which are pairwise orthogonal, meaning $e_F e_G = 0$ whenever $F \neq G$, and represent a linear basis of S_n over \mathbb{Z}_p .

The scheme is ring-based and involves a private plaintext ring P and a ciphertext ring C , where P is a retract of C , i.e., $P \subset C$ and at the same time P is a factor ring of C . Both P and C are rings of the above form. Let I be an ideal of C such that C/I is isomorphic to P .

The data owner, D , generates C and I during the key generation algorithm within the scheme. Parameters n, r , and p are chosen such that p is a large prime, $r > n$, and let $P = S_n$. The ciphertext ring $C = S_r$ is randomly generated from $P = S_n$ by $\{x_i\}_{i=n+1}^r$, where $r > n$. In other words, the ring $C = S_r$ expands upon the generators of $P = S_n$. The data owner then chooses an ideal I by randomly selecting elements

$$x_m - w_m(x_1, x_2, \dots, x_{m-1})$$

for $m = n + 1, \dots, r$, where $w_m(x_1, x_2, \dots, x_{m-1})$ represents an idempotent element of S_{m-1} . Observe that C/I is isomorphic to P by construction. The data owner next rewrites C in terms of its orthogonal basis $\{e_i\}_{i=1}^{2^r}$ and applies a random permutation π to the orthogonal set generators. D concludes the key generation algorithm with a final transformation between this permuted basis, given by

$$C = \langle e_1, e_2, \dots, e_{2^r} \mid p \cdot 1 = 0, e_i^2 = e_i, \text{ and } e_i e_j = 0 \text{ for all } i, j \rangle,$$

and a *triangular basis*. This triangular basis is initially constructed by letting

$$t_m = \sum_{i=1}^k e_k$$

for $1 \leq k \leq 2^r$. Next, m rows are randomly repeated in this basis and the final, public, ciphertext ring is given by

$$\tilde{C} = \langle t_1, t_2, \dots, t_{2^r+m} \mid p \cdot 1 = 0, e_i^2 = e_i, \text{ and } e_i e_j = 0 \text{ for all } i, j \rangle.$$

Note that in the published basis, homomorphic addition occurs component-wise and the homomorphic multiplication operation is given by $t_i t_j = t_i$ whenever $i \leq j$.

Elements in the orthogonal basis are not uniquely represented by elements in the triangular basis. For example, if $t_1 = e_1, t_2 = e_1, t_3 = e_1 + e_2$, and $t_4 = e_1 + e_2$, then the element e_2 in the orthogonal basis can be written in the triangular basis in four ways: $t_3 - t_1, t_3 - t_2, t_4 - t_1$, or $t_4 - t_2$. This is used to the key holder's advantage during encryption.

A plaintext element $x \in P$ is encrypted as

$$[[x]] = x + E_0$$

where

$$E_0 = \sum_{j=n+1}^r (x_j - w_j(x_1, \dots, x_{j-1})) \cdot h_j(x_1, \dots, x_r)$$

for random $h_j \in C$. The element E_0 therefore belongs to the ideal I . Covert $[[x]]$ from the standard basis to the orthogonal basis, then to the triangular basis. Because the orthogonal elements do not necessarily have a unique expression in the triangular basis, the triangular expression of each orthogonal basis element is randomly selected from possible representations during encryption.

To decrypt, the data owner simply converts the ciphertext back to the basis $\{x_i\}_{i=1}^r$, and then replaces x_j with $w_j(x_1, \dots, x_{j-1})$ for $j = n + 1, \dots, r$.

2. Privacy-preserving classification

The term *privacy-preserving classification* (PPC) broadly describes a

situation in which a user is able to classify her datapoint using a data owner's learned model, while neither party learns any information about the other party's data [14]. A major application of PPC lies in the medical field. With PPC, a patient can perform medical analyses of her medical data without revealing her personal information.

Differential privacy is one approach to which has seen some success with training various classification algorithms [31]. This method, however, is built upon having a large database and loses its utility when the goal is to classify a *single* datapoint [12]. Leveled homomorphic encryption schemes (LHE), a variant on fully homomorphic encryption which allow for polynomial computation up to a predefined depth, such as YASHE have been used in an application of neural networks to encrypted data called CryptoNets [12]. ML Confidential [1] used LHE to run classification using Linear Means and Fisher's Linear Discriminant classifiers.

In addition, there are methods which are not based entirely on fully homomorphic encryption. The authors in Ref. [14] construct efficient protocols for privacy-preserving classification via hyperplane detection, Naive Bayes, and decision trees using two additively homomorphic encryption schemes, Quadratic Reciprocity (QR) [32] and Paillier [33], and one leveled homomorphic encryption scheme, HELib [30].

2.1. Private naive bayes

The Naive Bayes classifier is based off of Bayes Theorem, which states that

$$\Pr(G = G_i|X) = \frac{\Pr(X|G = G_i)\Pr(G = G_i)}{\Pr(X)}$$

where $G = G_i$ denotes the possible classes for a data point X . In other words, we compute the posterior probability $\Pr(G = G_i|X)$ using prior knowledge combined with observed data.

The Naive Bayes model assumes that the inputs are conditionally independent in each class, hence the “naive” title. Naive Bayes is able to provide fast results which surprisingly are often better than more sophisticated methods. The Naive Bayes classifier does not consider correlation among features, thus lowering the variance while increasing the bias and hence avoiding over fitting on the training set. Due to this it has become a benchmark used by the community in data analysis.

The classification of a vector X with d features is given simply by taking the maximum posterior probability over all classes and applying Bayes' Theorem:

$$\operatorname{argmax}_{G_i \in G} \Pr(G = G_i) \prod_{j=1}^d \Pr(X = X_j|G = G_i).$$

Commonly in medical applications each attribute X_j of a vector X can take on only a discrete set of values. Then, all of the information needed to classify an arbitrary data point can be succinctly represented in a collection of vectors and matrices. This representation will enable us to collect all data needed in order to perform a classification. The ease with which all this data can be represented will be of great use in the protocol which follows. Let $P = (P_i)$ be a vector where

$$P_i = \Pr(G = G_i)$$

and let $T = (T_{i,j})$ be a matrix where

$$T_{i,j} = \Pr(X = X_j|G = G_i).$$

Given tables P and T , any new data point could be classified by looking up the probability for each attribute, given each class, and computing the argmax of the resulting values.

Bost et al. implement private naive bayes classification between a client and a service provider [14]. The client has a single data vector X which she would like to keep private. The service provider owns a private model w . Given a classification algorithm C the client should be able to learn $C(X, w)$, the classification of her vector X using the model

w , without learning any partial information about the model or giving away any information about her input.

Let square brackets $\llbracket a \rrbracket_{\mathcal{P}}$ denote the encryption of a value a under Paillier. The security of these schemes provides semantic security for the authors' protocols, and an honest-but-curious adversary model is used.

Assume there are r possible classes, d features per data point, and that there are a finite number of values each attribute can take. Let P denote the vector of class probabilities, where $P_i = \Pr(G_i)$, and let T denote the collection of tables given by $T_{i,j}(X) = \Pr(X = X_j|G = G_i)$. Assume each vector X has d features, and that there are c possible classes. The authors' protocol runs as follows:

- 1 The service provider encrypts tables P and T using Paillier.
- 2 The server sends the encrypted tables to the client.
- 3 The client computes

$$\llbracket P_i \rrbracket_{\mathcal{P}} = \llbracket P_i \rrbracket_{\mathcal{P}} \cdot \prod_{j=1}^d \llbracket T_{i,j}(x_j) \rrbracket_{\mathcal{P}}$$

For $i = 1, \dots, c$.

- 4 The client uses the server to compute $i = \operatorname{argmax}_i P_i$.
- 5 Client outputs i .

Note that Step 3 of the above protocol requires the encryption scheme to be multiplicatively homomorphic. Bost et al. implement the model using the logarithm of the probability distributions, where $p_i = \log(\Pr(G = G_i|X))$. Therefore, the posterior probabilities in Step 3 are computed using Paillier's additively homomorphic property. Computing the argmax in Step 4 requires an additively homomorphic scheme. Bost et al. use Paillier and QR, combined with an algorithm for changing the encryption scheme, in order to compute the argmax. In the next section, we describe our adapted version of the protocol which performs classification using a single encryption scheme.

3. Proposed method for fully homomorphic private classification with naive bayes

The model described below varies in several slight but important ways from the previous protocol. First, as stated previously, this protocol is designed for an encryption scheme which is both private-key and fully homomorphic. A fully homomorphic scheme adds flexibility to the computation by allowing the posterior probabilities to be computed using homomorphic multiplication or in the logarithmic model using homomorphic addition. Furthermore, the presented model assumes direct communication between the service provider and the data owner. In other words, there is not a trusted server acting as intermediary between the parties. The protocol could easily be adapted to include a cloud service provider to carry out computations, if desired.

The protocol is designed as follows. Assume that a Data Owner, Alice, wishes to classify her vector X which contains d features based off of a learned model w owned by a Classification Model Owner, Bob. The group of classes G contains r distinct classes, G_1, \dots, G_r . During this protocol Bob should learn no unnecessary information about the input provided by Alice, and Alice should learn nothing but the predicted class index of X .

Because Bob has already computed a learned model, he prepares two tables. First, Bob prepares table P represented as a column vector of degree r where $P_i = \Pr(G = G_i)$, the prior probability on class G_i . Next he prepares a table T , where entry T_{ij} represents $\Pr(X = X_j|G = G_i)$.

There is a certain amount of information which Alice *must* know in order to carry out the protocol. Namely, Alice must know that the data vector X has p features and that there are exactly r classes. However, she should not be able to deduce any information about the conditional class probabilities associated with q features, or the r class probabilities.

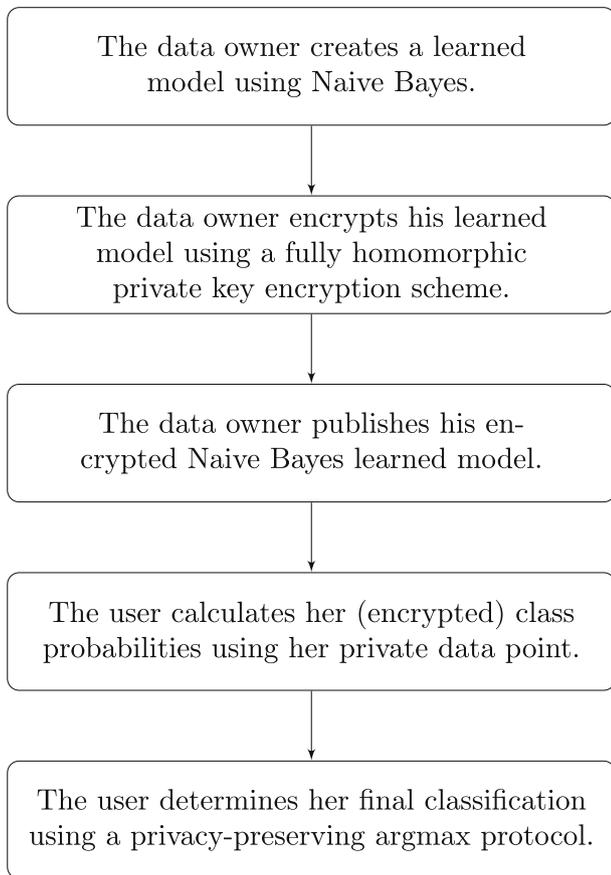


Fig. 1. Private naive bayes classification outline.

Private fully homomorphic Naive Bayes classification proceeds as follows. See Fig. 1 for an illustration of this protocol.

- 1 Bob prepares tables P and T and publishes their encryptions, $\llbracket P \rrbracket$ and $\llbracket T \rrbracket$.
- 2 For each class G_i for i from 1 to r , Alice computes

$$\begin{aligned} \llbracket \Pr(G_i|X) \rrbracket &= \llbracket \Pr(G_i) \rrbracket \cdot \prod_{j=1}^p \llbracket \Pr(X_j|G_i) \rrbracket \\ &= \llbracket P_i \rrbracket \cdot \prod_{j=1}^p \llbracket T_{ij} \rrbracket \\ &= \llbracket P_i \cdot \prod_{j=1}^p T_{ij} \rrbracket. \end{aligned}$$

3 Alice computes

$$i = \operatorname{argmax}_{1 \leq i \leq r} \llbracket \Pr(G_i|X) \rrbracket$$

using a private argmax protocol and outputs i .

If the model was designed using logarithms of the probabilities, then the multiplication operations in Step 2 should be replaced with addition.

There are two parties to consider when discussing the security of this protocol: the privacy of the Data Owner Alice's information as well as the privacy of the Classification Model Owner Bob's learned model. The privacy of the learned model is derived entirely from the security of the encryption scheme used. Discussion of the privacy of Alice's information, however, requires knowledge of the argmax protocol called in Step 3.

3.1. Private argmax

Previous approaches to private argmax using public-key encryption allow Alice to mask her value with random noise encrypted under Bob's public key. This approach will not work with private-key encryption because Alice cannot encrypt random noise.

The algorithm for computing argmax using a private-key FHE scheme is described via a series of improvements on an initial protocol which leaks information. Denote $\mathcal{E}_i = \llbracket P_i \rrbracket$ and the set of all \mathcal{E}_i as \mathcal{E} . First, suppose Alice computes argmax by sending the set of encrypted probability values \mathcal{E} to M . Then, Bob decrypts each value and sends Alice the index of the highest value using an asymmetric encryption protocol to hide the value of the index from an eavesdropper. While Alice has learned nothing about the encrypted model, Bob learns not only which class Alice's data point belongs to but also the exact probabilities for each class.

Suppose instead that Alice performs a permutation π on the class probabilities then sends $\pi(\mathcal{E})$ to Bob. Bob decrypts the values, determines which is largest, and sends that index to Alice, who reverses the permutation to determine her class. Note that it is not necessary to hide the value of the index sent to Alice from any eavesdroppers in this scenario, as the permutation π randomized the value. However, while this method prevents Bob from determining which probability is associated with each class specifically it does not prevent him from learning the class probabilities themselves.

Next, Alice attempts to hide the class probabilities by breaking down the computation of argmax to a series of comparisons. In particular, she randomly selects two encrypted probabilities $\mathcal{E}_i, \mathcal{E}_k \in \mathcal{E}$. She then sends the value

$$\mathcal{E}^* = \mathcal{E}_i - \mathcal{E}_k = \llbracket P_i - P_k \rrbracket$$

to Bob. Now when Bob decrypts \mathcal{E}^* he learns the value $P_i - P_k \in (-p/2, p/2]$, but not either of the probabilities themselves. If this value is negative, Bob sends the bit $b = 0$ to Alice, and sends $b = 1$ otherwise. Bob is then able to determine which value was the larger one based on the initial selection of P_i and P_k and rules out the smaller of the two values. The procedure is repeated until only one value remains.

Because the value sent to Bob is randomly ordered, publishing the bit 0 or 1 will also appear random to any observer hence does not compromise privacy. However, Bob recovers some partial information about Alice's data. He recovers a collection of r values representing the difference between pairs of the posterior probabilities for different $i \in [r]$. While permuting the elements keeps him from knowing which pairs of elements correspond to which values, it is not outside the realm of possibility that he could learn partial information about Alice's private data using this information.

Therefore, Alice needs to mask the value given by the difference between each pair of probabilities. To execute this, we take advantage of the fully homomorphic properties of our protocol by constructing a family \mathcal{F} of additively homomorphic, *monotone functions* that commute with encryption. A function $f: R \rightarrow R$ commutes with encryption if

$$f(\llbracket m \rrbracket) = \llbracket f(m) \rrbracket.$$

The additively homomorphic property of f guarantees that $f(m + n) = f(m) + f(n)$.

A privacy-preserving algorithm for argmax using a family of additively homomorphic monotone functions, \mathcal{F} , proceeds as follows:

- 1 Set $I = \{1, 2, \dots, r\}$.
- 2 **while** $|I| > 1$ **do**
- 3 A computes a random permutation π on I .
- 4 A randomly chooses $f \leftarrow \mathcal{F}$ and computes the values $f(\mathcal{E}_{\pi(1)})$ and $f(\mathcal{E}_{\pi(2)})$.
- 5 A uses the additive homomorphic properties of encryption and f as well as the commutative property of the function to evaluate

$$\begin{aligned} \mathcal{E}^* &= f(\mathcal{E}_{\pi(1)}) - f(\mathcal{E}_{\pi(2)}) \\ &= f(\llbracket P_{\pi(1)} \rrbracket) - f(\llbracket P_{\pi(2)} \rrbracket) \\ &= \llbracket f(P_{\pi(1)} - P_{\pi(2)}) \rrbracket \end{aligned}$$

6 A sends \mathcal{E}^* to B.

7 B decrypts \mathcal{E}^* and recovers

$$f(P_{\pi(1)} - P_{\pi(2)})$$

If this value is negative, B sends the bit $b = 0$ to A, otherwise send $b = 1$.

8 If $b = 0$, remove $\pi(1)$ from I . Otherwise remove $\pi(2)$.

9 end while

10 A returns I .

During this protocol, B (Bob) collects r values representing the result of a monotone function applied to the difference between random pairs of the posterior probabilities. The application of an unknown monotone function to this difference prevents him from learning partial information from the decrypted value.

3.2. Security

The proposed protocol is secure in the honest-but-curious setting. In this setting, all parties are assumed to take part in the protocol honestly. However, they will use any information they can obtain during an honest execution of the protocol in order to attempt to deduce further information. In this model, parties should not be able to learn any information beyond their protocol output.

During protocol execution, Alice only has access to encrypted values. The private argmax protocol allows Alice to determine which encrypted value is the largest out of a set of values, but does not give her access to the actual values. Therefore, the security of Bob's data is reliant upon the security of the fully homomorphic encryption scheme that is implemented.

The only information Bob receives during protocol execution is obtained during the private argmax protocol. Bob receives the value

$$f(P_{\pi(1)} - P_{\pi(2)})$$

during each elimination round of the argmax protocol. Because of the permutation π , Bob does not know which posterior probabilities are being compared, or, in the case of binary classification, in what order they are being compared. Furthermore, the function f , selected from a family \mathcal{F} of additively homomorphic, monotone functions that commute with encryption, hides the exact value of the difference between the randomized posterior probabilities from Bob. Therefore, at the end of protocol execution, Bob is unable to determine any ordering on the posterior probabilities Alice has computed, and cannot determine her final classification.

4. Implementation

The proposed protocol is implemented using the GKS encryption scheme, where coefficients in \mathbb{Z}_p are taken from $(-p/2, p/2]$ and correspond to positive and negative integers, and

$$\mathcal{F} = \{f: R \rightarrow R: f(m) = km\}.$$

Table 1

Unencrypted versus encrypted experimental results under 10×10 -fold cross validation.

	Time (s)	Accuracy	Sensitivity	Specificity	Precision	NPV
Unencrypted	0.00001	0.96003	0.93389	0.97410	0.95100	0.96476
Encrypted	0.40298	0.96003	0.93389	0.97410	0.95100	0.96476

for a randomly chosen integer k up to 20 bits. The implementation uses a 198-bit prime p .

Protocols were implemented in C++ on a MacBook Pro using El Capitan, a 2.3 GHz Intel Core i7, with 16 GB memory. The GNU Multiple Precision Library (GMP) [34] was used to allow for integer storage above the built-in data type limits in C++.

4.1. Encoding

Fully homomorphic computation can only be utilized over a fully homomorphic encoding. Gribov et al. provide a straightforward method of implementing a fully homomorphic encoding in the GKS scheme [15]. Any fully homomorphic embedding of \mathbb{Z}_p into the ring S_n will suffice. Specifically, it is sufficient to map the element 1 of \mathbb{Z}_p to any idempotent element of S .

Elements in the scheme have coefficients which lie in the integer field \mathbb{Z}_p . Therefore floating point numbers must be encoded as integers in \mathbb{Z}_p . In order to encode floating point values, the most straightforward approach is to scale the floating point values to integer values with a fixed precision. We can scale a floating point value to yield n levels of precision via an encoding function $\text{Encode}(x) = \lfloor x \cdot 10^n \rfloor$. Then, this encoding is extended to the ring S_n by mapping x to a fixed idempotent value in S and assigning random values in \mathbb{Z}_p to the remaining coefficients.

To decode, the coefficient of the fixed idempotent value in S is retrieved and multiplied by 10^{-n} . This encoding requires the user to keep track of the number of multiplications performed to properly decode after multiplication.

This encoding suffers from overflow over the prime modulus p if too many units of decimal precision are required. With each multiplication performed, the size of the input grows exponentially.

The model in our experiments was trained with five decimal points precision, and the unencrypted experimental results were obtained with five decimal points precision. In the encrypted experiments all values were initially encrypted with three decimal points precision. Table 1 shows that no loss in accuracy resulted from this change in decimal precision.

5. Results on breast cancer classification

Data from the UCI Machine Learning Repository was used to test the performance of the protocols [16]. Specifically, the Breast Cancer Wisconsin (Original) Data Set was used, which contains 683 complete data points each containing an ID along with 9 attributes and a binary classification.

The data gives measurements taken from fine-needle aspirate (FNA) biopsies of benign and malignant breast tumors. These nine attributes include clump thickness, uniformity of cell size, uniformity of cell shape, marginal adhesion, single epithelial cell size, bare nuclei, bland chromatin, normal nucleoli, and mitoses. A clinician measured and rated each attribute on a scale of 1–10 at the time it was collected. Lower values correspond to what clinicians expect to see in a benign case and higher values correspond to the malignant case. Previous research has found that while each measurement holds clinical significance in diagnosing a breast tumor as benign or as malignant, a single attribute is not enough to distinguish between the two cases [35].

A Naive Bayes learned model was created and encrypted using the GKS encryption scheme. We performed 10×10 -fold cross validation to

evaluate the performance of the learned model. Furthermore, because the data set is unbalanced with a higher proportion of benign tumors, random oversampling of the malignant tumor samples was performed while training the models. A positive case denotes a case where the tumor is diagnosed as malignant and a negative case refers to a tumor which is benign.

Additive smoothing was performed on the values in the tables of class and prior probabilities prior to testing to prevent error in the case of zero or near zero probabilities. Specifically, each probability was increased by 0.1, and any value which was greater than or equal to 1 after smoothing was reset to 0.999. The size of the ciphertext ring in the experiments was $2^8 + 50 = 306$, and the prime modulus p was 198 bits.

Classification was tested on data points not included in the training set in both encrypted and unencrypted formats and performed 10×10 -fold cross validation. Table 1 shows a comparison the average performance of the encrypted and unencrypted experiments. Results include classification time for a single data point in seconds, accuracy, sensitivity, precision, specificity, and the true negative rate.

The time increase between encrypted and unencrypted computation is expected and occurs in all current fully homomorphic encryption methods. For the example provided above, where a single user wishes to classify their data, classification in approximately half a second is well within a reasonable time range for medical applications. Bost et al. report a classification time of 0.479 s on the same data set, although it is not clear how large of a testing set was used to generate this score [14].

Other statistics we provide include Sensitivity, Specificity, Precision, and the Negative Predictive Value (NPV). In the case of breast cancer specifically it is crucial to positively identify all malignant tumors for timely medical intervention. Our results yield high precision and sensitivity, meaning there is a high rate of true positives and a low rate of false positives.

Specificity and the NPV are also known as Inverse Recall and Inverse Precision, as they provide similar information for negative classifications. Specificity describes the proportion of negative cases which were classified as negative and NPV describes the proportion of cases classified as negative which were negative in reality. Our results show both high Specificity and high NPV, pointing to a high rate of true negatives and a low rate of false negatives.

6. Discussion

Naive Bayes often outperforms more sophisticated classification methods for medical diagnosis, such as prediction of heart disease [4] and rheumatological diagnosis [5]. To compare the performance of the Naive Bayes classifier on the Wisconsin breast cancer data set, decision tree, support vector machine (SVM), k -nearest neighbors, and logistic regression classification were implemented in Python 3 under 10-fold cross validation [3]. Table 2 shows the results of each classifier, with the highest score for each metric highlighted in red.

Decision trees were limited to a maximum depth of 5, a minimum of 3 samples per split, and a minimum of 3 samples per leaf. Gini impurity was used to measure the quality of the split during training. Support vector machine (SVM) was implemented with a linear kernel. The k -nearest neighbors classifier was implemented using the ball tree algorithm with 5 neighbors, a leaf size of 30, and uniformly weighted

Table 2

Comparison to other machine learning classifiers, with the highest score for each metric in red.

	Accuracy	Sensitivity	Specificity	Precision	NPV
Naive Bayes	0.96003	0.93389	0.97410	0.95100	0.96476
Decision Tree	0.96628	0.96171	0.97500	0.98699	0.93487
Linear SVM	0.97069	0.96843	0.97500	0.98651	0.94500
k -Nearest Neighbors	0.97065	0.96838	0.97482	0.98646	0.94544
Logistic regression	0.96920	0.97298	0.96232	0.98017	0.95276

points, using the Euclidean distance measure. The logistic regression model was built using the LIBLINEAR algorithm for optimization of the learning function [36].

Naive Bayes outperforms the tested classifiers in terms of negative predicted value. In other words, the Naive Bayes classifier yields the highest probability that a patient with a negative diagnosis result does not have the disease in reality. The specificity of the Naive Bayes classifier differs from the highest achieved specificity by only 0.09%, meaning that if a person does not have a malignant tumor, they are likely to receive a negative result. While other classifiers outperform Naive Bayes on the other metrics, Naive Bayes remains competitive, differing by at most 4%.

The proposed method performs classification in an average of 0.40298 s, compared to 0.479 s per classification in the experiments of Bost et al. [14]. Adding in an assumption of 20 ms per one-way interaction due to network latency, the proposed method requires an average of 0.462 ms per classification. The computation time required depends largely on the experimental hardware. However, the presented method has a lower communication requirement – it requires 3 interactions between the data owner and model owner for classification on the breast cancer dataset, compared to 14 interactions using the method of Bost et al. Private classification methods for more sophisticated classifiers are even more time consuming. Wu et al. perform private decision tree classification on a tree with 12 decision nodes in approximately 0.545 s [37]. Bost et al. report average decision tree classification times on a 4 node tree in approximately 2.085 s [14]. Rahulathavan et al. perform private SVM classification on the same breast cancer data set in an average of 7.71 s [38].

The amount of time a clinician may expect this protocol to take will depend upon the number of classes and the number of features within the data set being analyzed. The computational bottleneck within fully homomorphic encryption schemes occurs during homomorphic multiplication operations. Testing the proposed implementation of the GKS scheme for encryption speed was carried out with a variety of sizes for the prime modulus p , where random noise was drawn from the range $[0, p/2)$. With a 32-bit modulus, homomorphic multiplication was computed in an average of 0.0229 s, and increased linearly to an average of 0.0402 s per multiplication with a 1024-bit prime modulus.

ML Confidential [1] implements a leveled homomorphic encryption scheme using a modified version of a scheme presented by Brakerski [39]. With a 128-bit prime modulus, the authors report an average multiplication time of 0.853 s. Halevi and Shoup's HELib [30], associated with IBM Research, implements the BGV leveled fully homomorphic encryption scheme [23] with the Smart-Vercauteren SIMD ciphertext packing techniques [40] and additional Gentry-Halevi-Smart optimizations [41]. Authors provide preliminary multiplication times ranging from 25.7 s to 473 s based on varying sizes of key generation parameters [42], although a March 2018 update to the project's GitHub repository announced forthcoming speedups. An implementation of HELib carried out by Khedr et al., in 2015 performed multiplication in an average of 3.6 s [43]. The authors in this work also propose an improvement on the Gentry-Sahai-Waters (GSW) scheme [26] within a framework they title Secure Homomorphic Implementation of Encrypted Data-Classifiers (SHIELD), with reported multiplication time of 0.372 s.

7. Conclusions

Private-key fully homomorphic encryption can classify medical data efficiently. Similar techniques could enable researchers and clinicians to utilize private medical data and models which they cannot access in the clear. Hospitals and companies with trained assisted diagnosis systems could use this technology to provide access to their models without giving away their parameters, and doctors can use this software without revealing their patients' information.

Acknowledgments

Delaram Kahrobaei and Vladimir Shpilrain are partially supported by ONR (Office of Naval Research) grant N00014-15-1-2164.

References

- [1] T. Graepel, K. Lauter, M. Naehrig, M.L. Confidential, *Machine Learning on Encrypted Data*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 1–21.
- [2] H.N.A. Pham, E. Triantaphyllou, The Impact of Overfitting and Overgeneralization on the Classification Accuracy in Data Mining, Springer US, Boston, MA, 2008, pp. 391–431.
- [3] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning*, second Edition, Springer Series in Statistics, Springer New York Inc., New York, NY, USA, 2009.
- [4] S. Jyoti, A. Ujma, S. Dipesh, S. Sunita, Predictive data mining for medical diagnosis: an overview of heart disease prediction 17 (8) 43–48.
- [5] I. Kononenko, *Machine Learning for Medical Diagnosis: History, State of the Art and Perspective* 23 (1) 89–109.
- [6] J.H. Cheon, M. Kim, K. Lauter, Homomorphic Computation of Edit Distance, Springer Berlin Heidelberg, Berlin, Heidelberg, 2015, pp. 194–212.
- [7] J.R. Troncoso-Pastoriza, S. Katzenbeisser, M. Celik, Privacy preserving error resilient dna searching through oblivious automata, ACM Conference on Computer and Communications Security (CCS), ACM Press, ACM Press, Alexandria, Virginia, USA, 2007, pp. 519–528.
- [8] E. Ayday, J.L. Raisaro, U. Hengartner, A. Molyneaux, J.-P. Hubaux, Privacy-preserving Processing of Raw Genomic Data, Springer Berlin Heidelberg, Berlin, Heidelberg, 2014, pp. 133–147.
- [9] B. Fons, K. Stefan, K. Klaus, T. Pim, Privacy-preserving matching of DNA profiles, *Cryptol. ePrint Arch.* 2008 (2008) 203.
- [10] K. Lauter, A. López-Alt, M. Naehrig, *Private Computation on Encrypted Genomic Data*, Springer International Publishing, Cham, 2015, pp. 3–27.
- [11] M. Kim, K. Lauter, Private genome analysis through homomorphic encryption, *Cryptology ePrint Archive*, Report 2015/965, 2015. <http://eprint.iacr.org/2015/965>.
- [12] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, J. Wernsing, *CryptoNets: applying neural networks to encrypted data with high throughput and accuracy*, PMLR (2016) 201–210.
- [13] R. Shokri, V. Shmatikov, Privacy-preserving deep learning, *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security, CCS '15*, ACM, New York, NY, USA, 2015, pp. 1310–1321.
- [14] R. Bost, R. Ada Popa, S. Tu, S. Goldwasser, *Machine Learning Classification over Encrypted Data*, Symposium on Network and Distributed System Security, NDSS, 2015.
- [15] A. Gribov, D. Kahrobaei, V. Shpilrain, Private-key fully homomorphic encryption in rings, *Groups, Complexity, Cryptology* 10.
- [16] M. Lichman, UCI machine learning repository, <http://archive.ics.uci.edu/ml>, (2013).
- [17] R. Rivest, L. Adleman, M. Dertouzos, On data banks and privacy homomorphisms, *Foundations of Secure Computation* 4 (11) (1978) 165–179.
- [18] C. Peikert, A decade of lattice cryptography, *Found. Trends Theor. Comput. Sci.* 10 (4) (2016) 283–424.
- [19] C. Gentry, Fully homomorphic encryption using ideal lattices, *Proceedings of the Forty-first Annual ACM Symposium on Theory of Computing, STOC '09*, ACM, 2009, pp. 169–178.
- [20] Z. Brakerski, V. Vaikuntanathan, Fully homomorphic encryption from ring-LWE and security for key dependent messages, *Proceedings of the 31st Annual Conference on Advances in Cryptology, CRYPTO'11*, Springer-Verlag, Berlin, Heidelberg, 2011, pp. 505–524.
- [21] Z. Brakerski, V. Vaikuntanathan, Efficient fully homomorphic encryption from (standard) LWE, *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS '11*, IEEE Computer Society, Washington, DC, USA, 2011, pp. 97–106.
- [22] M. v. Dijk, C. Gentry, S. Halevi, V. Vaikuntanathan, Fully homomorphic encryption over the integers, *Advances in Cryptology – EUROCRYPT 2010*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2010, pp. 24–43.
- [23] Z. Brakerski, V. Vaikuntanathan, C. Gentry, Fully homomorphic encryption without bootstrapping, *Innovations in Theoretical Computer Science*, 2012.
- [24] J.W. Bos, K. Lauter, J. Loftus, M. Naehrig, Improved Security for a Ring-based Fully Homomorphic Encryption Scheme, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 45–64.
- [25] J. Alperin-Sheriff, C. Peikert, Faster bootstrapping with polynomial error, *Advances in Cryptology – CRYPTO 2014*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2014, pp. 297–314.
- [26] C. Gentry, A. Sahai, B. Waters, Homomorphic encryption from learning with errors: conceptually-simpler, asymptotically-faster, attribute-based, CRYPTO, Springer, 2013, pp. 75–92.
- [27] C. Gentry, S. Halevi, C. Peikert, N.P. Smart, Ring switching in BGV-style homomorphic encryption, *Security and Cryptography for Networks*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2012, pp. 19–37.
- [28] C. Gentry, S. Halevi, N.P. Smart, Fully homomorphic encryption with polylog overhead, *Advances in Cryptology – EUROCRYPT 2012*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 2012, pp. 465–482.
- [29] C. Gentry, S. Halevi, N.P. Smart, Better Bootstrapping in Fully Homomorphic Encryption, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 1–16.
- [30] S. Halevi, HELib: an Implementation of homomorphic encryption, <https://github.com/shaih/HELlib>, (2013).
- [31] C. Dwork, *Differential Privacy*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 1–12.
- [32] S. Goldwasser, S. Micali, Probabilistic encryption & how to play mental poker keeping secret all partial information, *Proceedings of the Fourteenth Annual ACM Symposium on Theory of Computing, STOC '82* ACM, New York, NY, USA, 1982, pp. 365–377.
- [33] P. Paillier, Public-key cryptosystems based on composite degree residuosity classes, *Advances in Cryptology – EUROCRYPT '99*, Lecture Notes in Computer Science, Springer, Berlin, Heidelberg, 1999, pp. 223–238.
- [34] The GNU MP Bignum Library. URL <https://gmplib.org/>.
- [35] W.H. Wolberg, O.L. Mangasarian, Multisurface method of pattern separation for medical diagnosis applied to breast cytology, *Proceedings of the National Academy of Sciences, U.S.A.* 87 (1990) 9193–9196.
- [36] Liblinear – a library for large linear classification. URL <https://www.csie.ntu.edu.tw/~cjlin/liblinear>.
- [37] D.J. Wu, T. Feng, M. Naehrig, K. Lauter, Privately Evaluating Decision Trees and Random Forests, (2016), pp. 335–355 4.
- [38] Y. Rahulamathavan, R. C. Phan, S. Veluru, K. Cumanan, M. Rajarajan, Privacy-preserving Multi-class Support Vector Machine for Outsourcing the Data Classification in Cloud vol. 11 (5) 467–479.
- [39] Z. Brakerski, Fully homomorphic encryption without modulus switching from classical gapsvp, in: R. Safavi-Naini, R. Canetti (Eds.), *Advances in Cryptology – CRYPTO 2012*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 868–886.
- [40] N.P. Smart, F. Vercauteren, Fully homomorphic simd operations, *Des. Codes Cryptogr.* 71 (1) (2014) 57–81.
- [41] C. Gentry, S. Halevi, N.P. Smart, Homomorphic evaluation of the aes circuit, in: R. Safavi-Naini, R. Canetti (Eds.), *Advances in Cryptology – CRYPTO 2012*, Springer Berlin Heidelberg, 2012, pp. 850–867.
- [42] S. Halevi, V. Shoup, Algorithms in helib, *cryptology ePrint archive*, Report 2014/106, 2014. <http://eprint.iacr.org/2014/106>.
- [43] A. Khedr, G. Gulak, V. Vaikuntanathan, Shield: scalable homomorphic implementation of encrypted data-classifiers, *IEEE Trans. Comput.* 65 (9) (2016) 2848–2858.