

Machine Learning Classification over Encrypted Data

Raphaël Bost
DGA*

Raluca Ada Popa
MIT

Stephen Tu
MIT

Shafi Goldwasser
MIT

Abstract

Machine learning classification is used in numerous settings nowadays, such as medical or genomics predictions, spam detection, face recognition, and financial predictions. Due to privacy concerns, in some of these applications, it is important that the data and the classifier remain confidential.

In this work, we construct three major classification protocols that satisfy this privacy constraint: hyperplane decision, Naïve Bayes, and decision trees. We also enable these protocols to be combined with AdaBoost. At the basis of these constructions is a new library of building blocks for constructing classifiers securely; we demonstrate that this library can be used to construct other classifiers as well, such as a multiplexer and a face detection classifier.

We implemented and evaluated our library and classifiers. Our protocols are efficient, taking milliseconds to a few seconds to perform a classification when running on real medical datasets.

1 Introduction

Classifiers are an invaluable tool in many settings today, such as medical or genomics predictions, spam detection, face recognition, and finance. Many of these applications handle sensitive data [WGH12, SG11, SG13], so it is important that the data and the classifier remain private.

Consider the typical setup with supervised learning, depicted in Figure 1. Supervised learning algorithms consist of two phases: (i) the training phase during which the algorithm learns a model w from a set of labeled examples, and (ii) the classification phase that runs a classifier C over a previously unseen feature vector x , using the model w to output a prediction $C(x, w)$.

In applications that handle sensitive data, it is important that the feature vector x and the model w remain secret to one or some of the parties involved. Consider the example of a medical study or a hospital having a model built out of the private medical profiles of some patients; the model is sensitive because it can leak information about the patients, and its usage has to be HIPAA¹ compliant. A client wants to use the model to make a prediction about her health (e.g., if she is likely to contract a certain disease, or if she would be treated successfully at the hospital), but does not want to reveal her sensitive medical profile. Ideally, the hospital and the client run a protocol at the end of which the client learns one bit (“yes/no”), and neither party learns anything else about the other party’s input. A similar setting arises for a financial institution (e.g., an insurance company) holding a sensitive model, and a customer wanting to estimate rates or quality of service based on her personal information.

Throughout this paper, we refer to this goal shortly as *privacy-preserving classification*. Concretely, a client has a private input represented as a feature vector x , and the server has a private model w as input. The way the model w is obtained is independent of our protocols here. For example, the server could have computed the model w after running the training phase on plaintext data as usual. Only the classification needs to be privacy-preserving: the client should learn $C(x, w)$ but nothing else about the model w , while the server should not learn anything about the client’s input or the classification result.

In this work, we construct efficient privacy-preserving protocols for three of the most common classifiers: hyperplane decision, Naïve Bayes, and decision trees, as well as a more general classifier combining these using AdaBoost. These

*Direction Générale de l’Armement. Work done while visiting MIT CSAIL. The views and conclusions contained herein are those of the author and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the DGA or the French Government.

¹Health Insurance Portability and Accountability Act of 1996




Figure 1: Model overview. Each shaded box indicates private data that should be accessible to only one party: the dataset and the model to the server, and the input and prediction result to the client. Each straight non-dashed rectangle indicates an algorithm, single arrows indicate inputs to these algorithms, and double arrows indicate outputs.

Learning algorithm	Classifier
Perceptron	Hyperplane decision
Least squares	Hyperplane decision
Fischer linear discriminant	Hyperplane decision
Support vector machine	Hyperplane decision
Naive Bayes	Naïve Bayes
Decision trees (ID3/C4.5)	Decision trees

Table 1: Machine learning algorithms and their classifiers, defined in Section 3.1.

classifiers are widely used – even though there are many machine learning algorithms, most of them end up using one of these three classifiers, as described in Table 1.

While generic secure multi-party computation [Yao82, GMW87, HKoS⁺10, MNPS04, BDNP08] can implement any classifier in principle, due to their generality, such schemes are not efficient for common classifiers. As described in Section 10.5, on a small classification instance, these tools [HKoS⁺10, BDNP08] ran out of memory on a powerful machine with 256GB of RAM, and on an artificially simplified classification instance, they ran ≈ 500 times slower than our protocols ran on the non-simplified instance.

Most existing work in machine learning and privacy [LP00, DHC04, WY04, ZW05, BDMN05, VKC08, GLN13] focuses on preserving privacy during the *training phase*, and does not address classification. The few works on privacy-preserving classification either consider a weaker security setting in which the client learns the model [BLN13] or focus on specific classifiers (*e.g.*, face detectors [EFG⁺09, SSW10, AB06, AB07]) that are useful in limited situations.

Designing efficient privacy-preserving classification faces two main challenges. The first is that the computation performed over sensitive data by some classifiers is quite complex (*e.g.*, decision trees), making it hard to support efficiently. The second is providing a solution that is more generic than the three classifiers: constructing a separate solution for each classifier does not provide insight into how to combine these classifiers or how to construct other classifiers. Even though these are three of the most common classifiers, various settings use other classifiers or use a combination of these three classifiers (*e.g.*, AdaBoost). We address these challenges using two key techniques.

Our main technique is to identify *a set of core operations* over encrypted data that underlie many classification protocols. We found these operations to be comparison, argmax, and dot product. We use *efficient* protocols for each one of these, either by improving existing schemes (*e.g.*, for comparison) or by constructing new schemes (*e.g.*, for argmax).

Our second technique is to design these building blocks in a *composable* way, with regard to *both* functionality and security. To achieve this goal, we use a set of sub-techniques:

- The input and output of all our building blocks are data encrypted with additively homomorphic encryption. In addition, we provide a mechanism to switch from one encryption scheme to another. Intuitively, this enables a building block’s output to become the input of another building block;
- The API of these building blocks is flexible: even though each building block computes a fixed function, it allows a choice of which party provides the inputs to the protocol, which party obtains the output of the computation, and whether the output is encrypted or decrypted;
- The security of these protocols composes using modular sequential composition [Can98].

We emphasize that the contribution of our building blocks library goes beyond the classifiers we build in this paper: a user of the library can construct other privacy-preserving classifiers in a modular fashion. To demonstrate this point, we use our building blocks to construct a multiplexer and a classifier for face detection, as well as to combine our classifiers using AdaBoost.

We then use these building blocks to construct novel privacy-preserving protocols for the three common classifiers. Some of these incorporate additional techniques, such as an efficient evaluation of a decision tree with fully homomorphic encryption (FHE) based on a polynomial representation requiring only a small number of multiplications and based on SIMD FHE slots (see Section 7.2).

Our last contribution is an implementation and evaluation of our building blocks and classifiers. We evaluate our classifiers on real datasets with private data about breast cancer, credit card approval, audiology, and nursery data; our algorithms are efficient, running in milliseconds up to a few seconds, and consume a modest amount of bandwidth.

The rest of the paper is organized as follows. Section 2 describes related work, Section 3 provide the necessary machine learning and cryptographic background, Section 4 presents our building blocks, Sections 5–8 describe our classifiers, and Sections 9–10 present our implementation and evaluation results.

2 Related work

Our work is the first to provide efficient privacy-preserving protocols for a broad class of classifiers.

Secure two-party computation protocols for generic functions exist in theory [Yao82, GMW87, LP07, IPS08, LP09] and in practice [HKoS⁺10, MNPS04, BDNP08]. However, these rely on heavy cryptographic machinery, and applying them directly to our problem setting would be far too inefficient as exemplified in Section 10.5.

Previous work focusing on privacy-preserving machine learning can be broadly divided into two categories: (i) techniques for private training, and (ii) techniques for private classification (recall the distinction from Figure 1). Most existing work falls in the first category, which we discuss in Section 2.1. Our work falls in the second category, where there has been less work done, which we discuss in Section 2.2. We also mention work related to the building blocks we use in our protocols in Section 2.3.

It is worth mentioning that our work on privacy-preserving classification is incomparable to work on differential privacy in the machine learning community (see *e.g.* [CMS11]). Our work aims to protect the confidentiality of user data, whereas differential privacy seeks to bound the amount of statistical inference that can be performed on a particular individual.

2.1 Privacy-preserving training

Existing techniques have been developed for privacy-preserving *training* for Naïve Bayes [VKC08, WY04, ZW05], decision trees [BDMN05, LP00], linear discriminant classifiers [DHC04], and more general kernel methods [LLM06].

Grapel *et al.* [GLN13] show how to train several machine learning classifiers using a somewhat homomorphic encryption scheme. They focus on a few simple classifiers (*e.g.* the linear means classifier), and do not elaborate on more complex algorithms such as support vector machines. They also support private classification, but in a weaker security model where the client learns more about the model than just the final sign of the classification. Indeed, performing the final comparison with fully homomorphic encryption (FHE) is impractical, a difficulty we overcome with an interactive setting.

2.2 Privacy-preserving classification

Little work has been done to address the general problem of privacy-preserving classification in practice; previous work focuses on a weaker security setting (in which the client learns the model) and/or only supports specific classifiers.

In Bos *et al.* [BLN13], a third party can compute medical prediction functions over the encrypted data of a patient using fully homomorphic encryption. In their setting, everyone (including the patient) knows the predictive model, and their algorithm hides only the input of the patient from the cloud. Our protocols, on the other hand, also hide the model from the patient. Their algorithms cannot be applied to our setting because they leak more information than just the bit

of the prediction to the patient. Furthermore, our techniques are notably different; using FHE directly for our classifiers would result in significant overheads.

Barni *et al.* [BFK⁺09, BFL⁺09] construct secure evaluation of linear branching programs, which they use to implement a secure classifier of ECG signals. Their technique is based on finely-tuned garbled circuits. By comparison, our construction is not limited to branching programs (or decision trees), keeps the model private, and is twice as fast on branching programs.

Other works [EFG⁺09, SSW10, AB06, AB07] construct specific face recognition or detection classifiers. We focus on providing a set of *generic* classifiers and building blocks to construct more complex classifiers. In Section 10.1.2, we show how to construct a private face detection classifier using the modularity of our techniques.

2.3 Work related to our building blocks

Two of the basic components we use are private comparison and private computation of dot products. These subjects have been well studied previously; see [Yao82, DGK07, DGK09, Veu11, LT05, AB06, KSS09] for comparison techniques and [AD01, GLLM05, Kil05, AB06] for techniques to compute dot products. Section 4.1 discusses how we build on these tools.

3 Background and definitions

3.1 Classification in machine learning algorithms

The user’s input x is a vector of d elements $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ and is also called a feature vector. The standard problem in classification is to evaluate a classification function $C_w : \mathbb{R}^d \mapsto \{c_1, \dots, c_k\}$ that takes as input the feature vector x ; the output is $k^* = C_w(x) \in [1 \dots k]$, the class to which x corresponds, based on a model w . We now describe how three popular classifiers work on regular, unencrypted data. They differ in the model w and the function C_w . For more details, we refer the reader to [BN06].

Hyperplane decision-based classifiers. For this classifier, the model w consists of k vectors in \mathbb{R}^d ($w = \{w_i\}_{i=1}^k$). The classifier is (cf. [BN06]):

$$k^* = \operatorname{argmax}_{i \in [k]} \langle w_i, x \rangle. \quad (1)$$

We now explain how Eq. (1) captures many common machine learning algorithms. A hyperplane based classifier typically works with a hypothesis space \mathcal{H} equipped with an inner product $\langle \cdot, \cdot \rangle$. This classifier usually solves a binary classification problem ($k = 2$): given a user input x , x is classified in class c_2 if $\langle w, \phi(x) \rangle \geq 0$, otherwise it is labeled as part of class c_1 . Here, $\phi : \mathbb{R}^d \mapsto \mathcal{H}$ denotes the feature mapping from \mathbb{R}^d to \mathcal{H} [BN06]. In this work, we focus on the case when $\mathcal{H} = \mathbb{R}^d$ and note that a large class of infinite dimensional spaces can be approximated with a finite dimensional space (as in [RR07]), including the popular gaussian kernel (RBF). In this case, $\phi(x) = x$ or $\phi(x) = Px$ for a randomized projection matrix P chosen during training. Notice that Px consists solely of inner products; we will show how to support private evaluation of inner products later, so for simplicity we drop P from the discussion.

To extend such a classifier from 2 classes to k classes, we use one of the most common approaches, *one-versus-all*, where k different models $\{w_i\}_{i=1}^k$ are trained to discriminate each class from all the others. The decision rule is then given by (cf. [BN06]) to be Eq. (1). This framework is general enough to cover many common algorithms, such as support vector machines (SVMs), logistic regression, and least squares.

Naïve Bayes classifiers. For this classifier, the model w consists of various probabilities: the probability that each class c_i occurs, namely $\{p(C = c_i)\}_{i=1}^k$, and the probabilities that an element x_j of x occurs in a certain class c_i – or said otherwise, the probabilities of the j -th component x_j of x to be v when x belongs to category c_i – namely $\{\{p(X_j = v | C = c_i)\}_{v \in D_j}\}_{j=1}^d\}_{i=1}^k$ where D_j is X_j ’s domain². The classification function, using a *maximum a*

²Be careful to distinguish between X_j , the probabilistic event representing the values taken by the j -th feature of user’s input, and x_j , the actual value taken by the specific vector x .

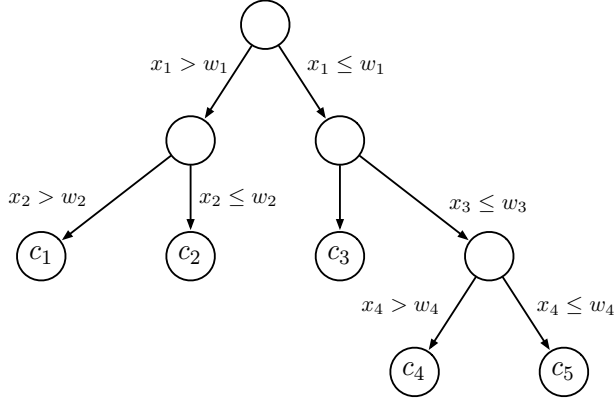


Figure 2: Decision tree

posteriori decision rule, works by choosing the class with the highest posterior probability:

$$\begin{aligned}
 k^* &= \operatorname{argmax}_{i \in [k]} p(C = c_i | X = x) \\
 &= \operatorname{argmax}_{i \in [k]} p(C = c_i, X = x) \\
 &= \operatorname{argmax}_{i \in [k]} p(C = c_i, X_1 = x_1, \dots, X_d = x_d)
 \end{aligned}$$

where the second equality follows from applying Bayes' rule.

The Naïve Bayes model assumes that $p(C = c_i, X = x)$ has the following factorization:

$$\begin{aligned}
 p(C = c_i, X_1 = x_1, \dots, X_d = x_d) \\
 &= p(C = c_i) \prod_{j=1}^d p(X_j = x_j | C = c_i),
 \end{aligned}$$

namely, each of the d features are conditionally independent given the class. For simplicity, we assume that the domain of the features values (the x_i 's) is discrete and finite, so the $p(X_j = x_j | C = c_i)$'s are probability masses.

Decision trees. A decision tree is a non-parametric classifier which works by partitioning the feature vector space one attribute at a time; interior nodes in the tree correspond to partitioning rules, and leaf nodes correspond to class labels. A feature vector x is classified by walking the tree starting from the root, using the partitioning rule at each node to decide which branch to take until a leaf node is encountered. The class at the leaf node is the result of the classification.

Figure 2 gives an example of a decision tree. The model consists of the structure of the tree and the decision criteria (in this case the thresholds w_1, \dots, w_4).

3.2 Cryptographic preliminaries

3.2.1 Cryptosystems

In this work, we use three additively homomorphic cryptosystems. A public-key encryption scheme HE is additively homomorphic if, given two encrypted messages $\text{HE.Enc}(a)$ and $\text{HE.Enc}(b)$, there exists a public-key operation \oplus such that $\text{HE.Enc}(a) \oplus \text{HE.Enc}(b)$ is an encryption of $a + b$. We emphasize that these are homomorphic *only for addition*, which makes them efficient, unlike fully homomorphic encryption [Gen09], which supports any function but at an expensive cost for multiplications. The cryptosystems we use are:

1. the QR (Quadratic Residuosity) cryptosystem of Goldwasser-Micali [GM82],

Type	Input A	Input B	Output A	Output B	Implementation
1	PK_P, PK_{QR}, a	SK_P, SK_{QR}, b	$[a < b]$	–	Sec. 4.1.1
2	$PK_P, SK_{QR}, \llbracket a \rrbracket, \llbracket b \rrbracket$	SK_P, PK_{QR}	–	$[a \leq b]$	Sec. 4.1.2
3	$PK_P, SK_{QR}, \llbracket a \rrbracket, \llbracket b \rrbracket$	SK_P, PK_{QR}	$a \leq b$	$[a \leq b]$	Sec. 4.1.2
4	$PK_P, PK_{QR}, \llbracket a \rrbracket, \llbracket b \rrbracket$	SK_P, SK_{QR}	$[a \leq b]$	–	Sec. 4.1.3
5	$PK_P, PK_{QR}, \llbracket a \rrbracket, \llbracket b \rrbracket$	SK_P, SK_{QR}	$[a \leq b]$	$a \leq b$	Sec. 4.1.3

Table 2: The API of our comparison protocol and its implementation. There are five types of comparisons each having a different setup.

2. the Paillier cryptosystem [Pai99], and
3. a leveled fully homomorphic encryption (FHE) scheme, HELib [Hal13]

3.2.2 Cryptographic assumptions

We prove that our protocols are secure based on the semantic security [Gol04] of the above cryptosystems. These cryptosystems rely on standard and well-studied computational assumptions: the Quadratic Residuosity assumption, the Decisional Composite Residuosity assumption, and the Ring Learning With Error (RLWE) assumption.

3.2.3 Adversarial model

We prove security of our protocols using the secure two-party computation framework for passive adversaries (or honest-but-curious [Gol04]) defined in Appendix B.1. To enable us to compose various protocols into a bigger protocol in a secure way, we invoke modular sequential composition (see Appendix B.2).

3.3 Notation

All our protocols are between two parties: parties A and B for our building blocks and parties C (client) and S (server) for our classifiers.

Inputs and outputs of our building blocks are either unencrypted or encrypted with an additively homomorphic encryption scheme. We use the following notation. The plaintext space of QR is \mathbb{F}_2 (bits), and we denote by $[b]$ a bit b encrypted under QR; the plaintext space of Paillier is \mathbb{Z}_N where N is the public modulus of Paillier, and we denote by $\llbracket m \rrbracket$ an integer m encrypted under Paillier. The plaintext space of the FHE scheme is \mathbb{F}_2 . We denote by SK_P and PK_P , a secret and a public key for Paillier, respectively. Also, we denote by SK_{QR} and PK_{QR} , a secret and a public key for QR.

For a constant b , $a \leftarrow b$ means that a is assigned the value of b . For a distribution \mathcal{D} , $a \leftarrow \mathcal{D}$ means that a gets a sample from \mathcal{D} .

4 Building blocks

In this section, we develop a library of building blocks, which we later use to build our classifiers. We designed this library to also be useful in constructing other classifiers than the ones described in our paper. The building blocks in this section combine existing techniques with either new techniques or new optimizations.

4.1 Comparison

We now describe our comparison protocol. In order for this protocol to be used in a wide range of classifiers, its setup needs to be *flexible*: namely, it has to support a range of choices regarding which party gets the input, which party gets the output, and whether the input or output are encrypted or not. Table 2 shows the various ways our comparison protocol can be used. In each case, each party learns *nothing else* about the other party’s input other than what Table 2 indicates as the output.

There are at least two approaches to performing comparison efficiently: using specialized homomorphic encryption, or using garbled circuits. For the former, [DGK07, DGK09] and [EFG⁺09] implement comparison with unencrypted inputs and Veugen [Veu11] performs comparison with encrypted inputs. To implement comparison using the second method, one could use the state-of-the-art garbling scheme of Bellare et al. [BHKR13], the short circuit for comparison of Kolesnikov et al. [KSS09] and a recent oblivious transfer scheme due to Asharov et al. [ALSZ13]. Since each garbled circuit can be used at most once, during the classification protocol, the parties need to both generate a fresh garbling and evaluate this garbled circuit. To compose with other building blocks that use homomorphic encryption, one also needs to convert from homomorphic encryption to garbled inputs and from a garbled output to homomorphic encryption. This task can be done using random shares as discussed in [KSS]. Both of these approaches to comparison yield good performance for our setting, and we preferred the first approach due to its simplicity.

We implemented each row of Table 2 by modifying existing protocols. We explain only the modifications here, and defer full protocol descriptions to Appendix A and proofs of security to Appendix C.1

4.1.1 Comparison with unencrypted inputs (Row 1)

We use a variation of the protocol described in [DGK07, DGK09, EFG⁺09], which we call DGK.

DGK. In this protocol, two parties A and B want to compare two private integers a (belonging to A) and b (belonging to B). B has a Paillier secret key SK_P and a QR secret key SK_{QR} , whereas A has the associated public keys (PK_P and PK_{QR}). At the end of the protocol, A outputs the encrypted bit $[t]$ where $t = (a < b)$. We refer the reader to [EFG⁺09] (c.f. Section 5) and [DGK07] (c.f. Section 3) for the full description and proof. (Veugen [Veu11] provides an alternative comparison protocol, called LSIC. LSIC performs less modular multiplications per invocation than DGK. However, our experimental evaluation shows that the number of rounds in LSIC – linear in the size of the input – is too slow for practical uses, so we use only DGK.)

4.1.2 Comparison with encrypted inputs (Rows 2, 3)

To develop our protocols, we require the ability to compare two encrypted inputs. More specifically, suppose that party A wants to compare two encrypted integers a and b , but party B holds the decryption key. To implement this task, we modify Veugen’s [Veu11] protocol to use DGK instead of LSIC as a subprocedure. This yields a protocol for the setup in Row 2. To ensure that A receives the plaintext output as in Row 3, B sends the encrypted result to A who decrypts it. Appendix A provides the detailed protocol.

4.1.3 Reversed comparison over encrypted data (Row 4, 5)

In some cases, we want the result of the comparison to be held by the party that does not hold the encrypted data. For this, we modify Veugen’s protocol to reverse the outputs of party A and party B : we achieve this by exchanging the role of party A and party B in the second part of the protocol, after invoking the DGK protocol. We do not present the details in the paper body because they are not insightful, and instead include them in Appendix A. This results in a protocol whose specification is in Row 4. To obtain Row 5, A sends the encrypted result to B who can decrypt it.

4.1.4 Negative integers comparison and sign determination

Negative numbers are handled by the protocols above unchanged. Even though the Paillier plaintext size is “positive”, a negative number simply becomes a large number in the plaintext space due to cyclicity of the space. As long as the values encrypted are within a preset interval $(-2^\ell, 2^\ell)$ for some fixed ℓ , Veugen’s protocol and the above protocols work correctly.

In some cases, we need to compute the sign of an encrypted integer $[[b]]$. In this case, we simply compare encryption of 0.

4.2 argmax over encrypted data

In this scenario, party A has k values a_1, \dots, a_k encrypted under party B 's secret key and wants party B to know the argmax over these values (the index of the largest value), but neither party should learn anything else. For example, if A has values $\llbracket 1 \rrbracket$, $\llbracket 100 \rrbracket$ and $\llbracket 2 \rrbracket$, B should learn that the second is the largest value, but learn nothing else. In particular, B should not learn the order relations between the a_i 's.

Our protocol for argmax is shown in Protocol 1. We now provide intuition into the protocol and its security.

Intuition. Let's start with a strawman. To prevent B from learning the order of the k values $\{a_i\}_{i=1}^k$, A applies a random permutation π . The i -th element becomes $\llbracket a'_i \rrbracket = \llbracket a_{\pi(i)} \rrbracket$ instead of $\llbracket a_i \rrbracket$.

Now, A and B compare the first two values $\llbracket a'_1 \rrbracket$ and $\llbracket a'_2 \rrbracket$ using the comparison protocol from row 4 of Table 2. B learns the index, m , of the larger value, and tells A to compare $\llbracket a'_m \rrbracket$ to $\llbracket a'_3 \rrbracket$ next. After iterating in this manner through all the k values, B determines the index m of the largest value. A can then compute $\pi^{-1}(m)$ which represents the argmax in the original, unpermuted order.

Since A applied a random permutation π , B does not learn the ordering of the values. The problem, though, is that A learns this ordering because, at every iteration, A knows the value of m up to that step and π . One way to fix this problem is for B to compare every pair of inputs from A , but this would result in a quadratic number of comparisons, which is too slow.

Instead, our protocol preserves the linear number of comparisons from above. The idea is that, at each iteration, once B determines which is the maximum of the two values compared, B should *randomize* the encryption of this maximum in such a way that A cannot link this value to one of the values compared. B uses the Refresh procedure for the randomization of Paillier ciphertexts. In the case where the "refresher" knows the secret key, this can be seen as a decryption followed by a re-encryption. If not, it can be seen as a multiplication by an encryption of 0.

A difficulty is that, to randomize the encryption of the maximum $\llbracket a'_m \rrbracket$, B needs to get this encryption – however, B must not receive this encryption because B has the key SK_P to decrypt it, which violates privacy. Instead, the idea is for A itself to add noise r_i and s_i to $\llbracket a'_m \rrbracket$, so decryption at B yields random values, then B refreshes the ciphertext, and then A removes the randomness r_i and s_i it added.

In the end, our protocols does $(k-1)$ encrypted comparisons of l bits integers and $7(k-1)$ homomorphic operations (refreshes, multiplications and subtractions). In terms of round trips, we have the ones due to the comparisons, to which we add one round trip per loop iteration *i.e.* $k-1$ round trips.

Proposition 4.1. *Protocol 1 is correct and secure in the honest-but-curious model.*

Proof intuition. The correctness property is straightforward. Let's argue security. A does not learn intermediary results in the computation because of the security of the comparison protocol and because she gets a fresh ciphertext from B which A cannot couple to a previously seen ciphertext. B does learn the result of each comparison – however, since A applied a random permutation before the comparison, B learns no useful information. See Appendix C for a complete proof.

4.3 Changing the encryption scheme

To enable us to compose various building blocks, we developed a protocol for converting ciphertexts from one encryption scheme to another while maintaining the underlying plaintexts. We first present a protocol that switches between two encryption schemes with the *same* plaintext size (such as QR and FHE over bits), and then present a different protocol for switching from QR to Paillier.

Concretely, consider two additively homomorphic encryption schemes E_1 and E_2 , both semantically secure with the same plaintext space M . Let $\llbracket \cdot \rrbracket_1$ be an encryption using E_1 and $\llbracket \cdot \rrbracket_2$ an encryption using E_2 . Consider that party B has the secret keys SK_1 and SK_2 for both schemes and A has the corresponding public keys PK_1 and PK_2 . Party A also has a value encrypted with PK_1 , $\llbracket c \rrbracket_1$. Our protocol, protocol 2, enables A to obtain an encryption of c under E_2 , $\llbracket c \rrbracket_2$ without revealing anything to B about c .

Proof Intuition. The idea is for A to add a random noise r to the ciphertext using the homomorphic property of E_1 . Then B decrypts the resulting value with E_1 (obtaining $x + r \in M$) and encrypts it with E_2 , sends the result to A which removes the randomness r using the homomorphic property of E_2 . Even though B was able to decrypt $\llbracket c \rrbracket_1$, B obtains $x + r \in M$ which hides x in an information-theoretic way (it is a one-time pad).

Protocol 1 argmax over encrypted data

Input A: k encrypted integers ($\llbracket a_1 \rrbracket, \dots, \llbracket a_k \rrbracket$), the bit length l of the a_i , and public keys PK_{QR} and PK_P

Input B: Secret keys SK_P and SK_{QR} , the bit length l

Output A: $\text{argmax}_i a_i$

```
1: A: chooses a random permutation  $\pi$  over  $\{1, \dots, k\}$ 
2: A:  $\llbracket \text{max} \rrbracket \leftarrow \llbracket a_{\pi(1)} \rrbracket$ 
3: B:  $m \leftarrow 1$ 
4: for  $i = 2$  to  $k$  do
5:   Using the comparison protocol (Sec. 4.1.3), B gets the bit  $b_i = (\text{max} \leq a_{\pi(i)})$ 
6:   A picks two random integers  $r_i, s_i \leftarrow (0, 2^{\lambda+l}) \cap \mathbb{Z}$ 
7:   A:  $\llbracket m'_i \rrbracket \leftarrow \llbracket \text{max} \rrbracket \cdot \llbracket r_i \rrbracket$   $\triangleright m'_i = \text{max} + r_i$ 
8:   A:  $\llbracket a'_i \rrbracket \leftarrow \llbracket a_{\pi(i)} \rrbracket \cdot \llbracket s_i \rrbracket$   $\triangleright a'_i = a_{\pi(i)} + s_i$ 
9:   A sends  $\llbracket m'_i \rrbracket$  and  $\llbracket a'_i \rrbracket$  to B
10:  if  $b_i$  is true then
11:    B:  $m \leftarrow i$ 
12:    B:  $\llbracket v_i \rrbracket \leftarrow \text{Refresh}[\llbracket a'_i \rrbracket]$   $\triangleright v_i = a'_i$ 
13:  else
14:    B:  $\llbracket v_i \rrbracket \leftarrow \text{Refresh}[\llbracket m'_i \rrbracket]$   $\triangleright v_i = m'_i$ 
15:  end if
16:  B sends to A  $\llbracket v_i \rrbracket$ 
17:  B sends to A the couple  $(\llbracket x_i \rrbracket, \llbracket y_i \rrbracket) = (\llbracket \bar{b}_i \rrbracket, \llbracket b_i \rrbracket)$ 
18:  A:  $\llbracket \text{max} \rrbracket \leftarrow \llbracket v_i \rrbracket \cdot \llbracket x_i \rrbracket^{-r_i} \cdot \llbracket y_i \rrbracket^{-s_i}$ 
19:  $\triangleright \text{max} = v_i - x_i \cdot r_i - y_i \cdot t_i$ 
20: end for
21: B sends  $m$  to A
22: A outputs  $\pi^{-1}(m)$ 
```

Note that, for some schemes, the plaintext space M depends on the secret keys. In this case, we must be sure that party A can still choose uniformly elements of M without knowing it. For example, for Paillier, $M = \mathbb{Z}_N^* \simeq \mathbb{Z}_p^* \times \mathbb{Z}_q^*$ where p and q are the private primes. However, in this case, A can sample noise in \mathbb{Z}_N that will not be in \mathbb{Z}_N^* with negligible probability $(1 - \frac{1}{p})(1 - \frac{1}{q}) \approx 1 - \frac{2}{\sqrt{N}}$ (remember N is large – 1024 bits in our instantiation).

Proposition 4.2. *Protocol 2 is secure in the honest-but-curious model.*

In our classifiers, we use this protocol for $M = \{0, 1\}$ and the encryption schemes are QR (for E_1) and an FHE scheme over bits (for E_2). In some cases, we might also want to switch from QR to Paillier (*e.g.* reuse the encrypted result of a comparison in a homomorphic computation), which has a different message space. Note that we can *simulate* the homomorphic XOR operation and a message space $M = \{0, 1\}$ with Paillier: we can easily compute the encryption of $b_1 \oplus b_2$ under Paillier when at most one of the b_i is encrypted (which we explain in the next subsection). This is the case in our setting because party A has the randomness r in the clear.

4.3.1 XOR with Paillier.

Suppose a party gets the bit b_1 encrypted under Paillier's encryption scheme, and that this party only has the public key. This party knows the bit b_2 in the clear and wants to compute the encryption of $\llbracket b_1 \oplus b_2 \rrbracket$.

To do so, we just have to notice that

$$b_1 \oplus b_2 = \begin{cases} b_1 & \text{if } b_2 = 0 \\ 1 - b_1 & \text{if } b_2 = 1 \end{cases}$$

Protocol 2 Changing the encryption scheme

Input A: $\llbracket c \rrbracket_1$ and public keys PK_1 and PK_2 **Input B:** Secret keys SK_1 and SK_2 **Output A:** $\llbracket c \rrbracket_2$

- 1: A uniformly picks $r \leftarrow M$
 - 2: A sends $\llbracket c' \rrbracket_1 \leftarrow \llbracket c \rrbracket_1 \cdot \llbracket r \rrbracket_1$ to B
 - 3: B decrypts c' and re-encrypts with E_2
 - 4: B sends $\llbracket c' \rrbracket_2$ to A
 - 5: A: $\llbracket c \rrbracket_2 = \llbracket c' \rrbracket_2 \cdot \llbracket r \rrbracket_2^{-1}$
 - 6: A outputs $\llbracket c \rrbracket_2$
-

Hence, it is very easy to compute an encryption of $b_1 \oplus b_2$ if we know the modulus N and the generator g (cf. Paillier's scheme construction):

$$\llbracket b_1 \oplus b_2 \rrbracket = \begin{cases} \llbracket b_1 \rrbracket & \text{if } b_2 = 0 \\ g[\llbracket b_1 \rrbracket]^{-1} \bmod N^2 & \text{if } b_2 = 1 \end{cases}$$

If we want to unveil the result to an adversary who knows the original encryption of b_1 (but not the secret key), we have to refresh the result of the previous function to ensure semantic security.

4.4 Computing dot products

For completeness, we include a straightforward algorithm for computing dot products of two vectors, which relies on Paillier's homomorphic property.

Protocol 3 Private dot product

Input A: $x = (x_1, \dots, x_d) \in \mathbb{Z}^d$, public key PK_P **Input B:** $y = (y_1, \dots, y_d) \in \mathbb{Z}^d$, secret key SK_P **Output A:** $\llbracket \langle x, y \rangle \rrbracket$

- 1: B encrypts y_1, \dots, y_d and sends the encryptions $\llbracket y_i \rrbracket$ to A
 - 2: A computes $\llbracket v \rrbracket = \prod_i \llbracket y_i \rrbracket^{x_i} \bmod N^2$ $\triangleright v = \sum y_i x_i$
 - 3: A re-randomizes and outputs $\llbracket v \rrbracket$
-

Proposition 4.3. *Protocol 3 is secure in the honest-but-curious model.*

4.5 Dealing with floating point numbers

Although all our protocols manipulate integers, classifiers usually use floating point numbers. Hence, when developing classifiers with our protocol library, we must adapt our protocols accordingly.

Fortunately, most of the operations involved are either additions or multiplications. As a consequence, a simple solution is to multiply each floating point value by a constant K (e.g. $K = 2^{52}$ for IEEE 754 doubles) and thus support finite precision. We must also consider the bit length for the comparisons. We show an example of a full analysis in Section 6 for the Naïve Bayes classifier.

5 Private hyperplane decision

Recall from Section 3.1 that this classifier computes

$$k^* = \operatorname{argmax}_{i \in [k]} \langle w_i, x \rangle.$$

Now that we constructed our library of building blocks, it is straightforward to implement this classifier securely: the client computes the encryption of $\llbracket \langle w_i, x \rangle \rrbracket$ for all $i \in [k]$ using the dot product protocol and then applies the argmax protocol (Protocol 1) to the encrypted dot products.

Protocol 4 Private hyperplane decision

Client's (C) Input: $x = (x_1, \dots, x_d) \in \mathbb{Z}^d$, public keys PK_P and PK_{QR}

Server's (S) Input: $\{w_i\}_{i=1}^k$ where $\forall i \in [k], w_i \in \mathbb{Z}^n$, secret keys SK_P and SK_{QR}

Client's Output: $\operatorname{argmax}_{i \in [k]} \langle w_i, x \rangle$

1: **for** $i = 1$ **to** k **do**

2: C and S run Protocol 3 for private dot product where C is party A with input x and S is party B with input w_i .

3: C gets $\llbracket v_i \rrbracket$ the result of the protocol.

$$\triangleright v_i \leftarrow \langle x, w_i \rangle$$

4: **end for**

5: C and S run Protocol 1 for argmax where C is the A, and S the B, and $\llbracket v_1 \rrbracket, \dots, \llbracket v_k \rrbracket$ the input ciphertexts. C gets the result i_0 of the protocol.

$$\triangleright i_0 \leftarrow \operatorname{argmax}_{i \in [k]} v_i$$

6: C outputs i_0

Proposition 5.1. *Protocol 4 is secure in the honest-but-curious model.*

6 Secure Naïve Bayes classifier

Section 3.1 describes the Naïve Bayes classifier. The goal is for the client to learn k^* without learning anything about the probabilities that constitute the model, and the server should learn nothing about x . Recall that the features values domain is discrete and finite.

As is typically done for numerical stability reasons, we work with the logarithm of the probability distributions:

$$\begin{aligned} k^* &= \operatorname{argmax}_{i \in [k]} \log p(C = c_i | X = x) \\ &= \operatorname{argmax}_{i \in [k]} \left\{ \log p(C = c_i) + \sum_{j=1}^d \log p(X_j = x_j | C = c_i) \right\} \end{aligned} \quad (2)$$

6.1 Preparing the model

Since the Paillier encryption scheme works with integers, we convert each log of a probability from above to an integer by multiplying it with a large number K (recall that the plaintext space of Paillier is large $\approx 2^{1024}$ thus allowing for a large K), thus still maintaining high accuracy. The issues due to using integers for bayesian classification have been previously studied in [TRMP12], even though their setting was even more restricting than ours. However, they use a similar idea to ours: *shifting* the probabilities logarithms and use fixed point representation.

As the only operations used in the classification step are additions and comparisons (*cf.* Equation (2)), we can just multiply the conditional probabilities $p(x_j | c_i)$ by a constant K so to get integers everywhere, while keeping the same classification result.

For example, if we are able to compute the conditional probabilities using IEEE 754 double precision floating point numbers, with 52 bits of precision, then we can represent every probability p as

$$p = m \cdot 2^e$$

where m binary representation is $(m)_2 = 1.d$ and d is a 52 bits integer. Hence we have $1 \leq m < 2$ and we can rewrite m as

$$m = \frac{m'}{2^{52}} \text{ with } m' \in \mathbb{N} \cap [2^{52}, 2^{53})$$

We are using this representation to find a constant K such that $K \cdot v_i \in \mathbb{N}$ for all i . As seen before, we can write the v_i 's as

$$v_i = m'_i \cdot 2^{e_i - 52}$$

Let $e^* = \min_i e_i$, and $\delta_i = e_i - e^* \geq 0$. Then,

$$v_i = m'_i \cdot 2^{\delta_i} \cdot 2^{e^* - 52}$$

So let $K = 2^{52 - e^*}$. We have $K \cdot v_i = m'_i \cdot 2^{\delta_i} \in \mathbb{N}$. An important thing to notice is that the v_i 's can be very large integers (due to δ_i), and this might cause overflows errors. However, remember that we are doing all this to store logarithms of probabilities in Paillier cyphertexts, and as Paillier plaintext space is *very* large (more than 1024 bits in our setting) and δ_i 's remain small³. Also notice that this *shifting* procedure can be done without any loss of precision as we can directly work with the bit representation of the floating points numbers.

Finally, we must also ensure that we do not overflow Paillier's message space when doing all the operations (homomorphic additions, comparisons, ...). If – as before – d is the number of features, the maximum number of bits when doing the computations will be $l_{max} = d + 1 + (52 + \delta^*)$ where $\delta^* = \max \delta_i$: we have to add the probabilities for the d features and the probability of the class label (the $d + 1$ term), and each probability is encoded using $(52 + \delta^*)$ bits. Hence, the value l used for the comparison protocols must be chosen larger than l_{max} .

Hence, we must ensure that $\log_2 N > l_{max} + 1 + \lambda$ where λ is the security parameter and N is the modulus for Paillier's cryptosystem plaintext space (cf. Section 4.1.2). This condition is easily fulfilled as, for a good level of security, we have to take $\log_2 N \geq 1024$ and we usually take $\lambda \approx 100$.

Let D_j be the domain of possible values of x_j (the j -th attribute of the feature vector x). The server prepares $kd + 1$ tables as part of the model, where K is computed as described just before:

- One table for the priors on the classes P : $P(i) = \lceil K \log p(C = c_i) \rceil$.
- One table per feature j per class i , $T_{i,j}$: $T_{i,j}(v) \approx \lceil K \log p(X_j = v | C = c_i) \rceil$, for all $v \in D_j$.

The tables remain small: P has one entry by category *i.e.* k , and T has one entry by category and feature value *i.e.* $k \cdot D$ where $D = \sum |D_j|$. In our examples, this represents less than 3600 entries. Moreover, this preparation step can be done once and for all at server startup, and is hence amortized.

6.2 Protocol

Let us begin with some intuition. The server encrypts each entry in these tables with Paillier and gives the resulting encryption (the encrypted model) to the client. For every class c_i , the client uses Paillier's additive homomorphism to compute $\llbracket p_i \rrbracket = \llbracket P(i) \rrbracket \prod_{j=1}^d \llbracket T_{i,j}(x_j) \rrbracket$. Finally, the client runs the argmax protocol, Protocol 1, to get $\text{argmax } p_i$. For completeness, the protocol is shown in Protocol 5.

Proposition 6.1. *Protocol 5 is secure in the honest-but-curious model.*

Proof intuition. Given the security property of the argmax protocol, Protocol 1, and the semantic security of the Paillier cryptosystem, the security of this classifier follows trivially, by invoking a modular composition theorem.

Efficiency. Note that the tables P and $\{T_{i,j}\}_{1 \leq i \leq k, 1 \leq j \leq d}$ can be prepared in advance. Hence the cost of constructing the tables can be amortized over many uses. To compute the encrypted probabilities p_i 's, the client runs d homomorphic

³If the biggest δ_i is 10, the ratio between the smallest and the biggest probability is of order $2^{2^{10}} = 2^{1024}$...

Protocol 5 Naïve Bayes Classifier

Client's (C) Input: $x = (x_1, \dots, x_d) \in \mathbb{Z}^d$, public key PK_P , secret key SK_{QR}

Server's (S) Input: The secret key SK_P , public key PK_{QR} and probability tables $\{\log p(C = c_i)\}_{1 \leq i \leq k}$ and $\left\{ \left\{ \log p(X_j = v | C = c_i) \right\}_{v \in D_j} \right\}_{1 \leq j \leq d, 1 \leq i \leq k}$

Client's Output: i_0 such that $p(x, c_{i_0})$ is maximum

- 1: The server prepares the tables P and $\{T_{i,j}\}_{1 \leq i \leq k, 1 \leq j \leq d}$ and encrypts their entries using Paillier.
 - 2: The server sends $\llbracket P \rrbracket$ and $\{\llbracket T_{i,j} \rrbracket\}_{i,j}$ to the client.
 - 3: For all $1 \leq i \leq k$, the client computes $\llbracket p_i \rrbracket = \llbracket P(i) \rrbracket \prod_{j=1}^d \llbracket T_{i,j}(x_j) \rrbracket$.
 - 4: The client runs the argmax protocol (Protocol 1) with the server and gets $i_0 = \text{argmax}_i p_i$
 - 5: C outputs i_0
-

operations (here multiplications) for each i , hence doing kd modular multiplications. Then the parties run a single argmax protocol *i.e.* $k - 1$ comparisons and $O(k)$ homomorphic operations. Thus, compared to non-encrypted computation, the overhead comes only from the use of homomorphic encryption operations instead of plaintext operations. Regarding the number of round trips, these are due to the argmax protocol: $k - 1$ runs of the comparison protocol and $k - 1$ additional roundtrips.

7 Private decision trees

A private decision tree classifier allows the server to traverse a binary decision tree using the client's input x such that the server does not learn the input x , and the client does not learn the structure of the tree and the thresholds at each node. A challenge is that, in particular, the client should not learn the path in the tree that corresponds to x – the position of the path in the tree and the length of the path leaks information about the model. The outcome of the classification does not necessarily leak the path in the tree

The idea is to express the decision tree as a polynomial P whose output is the result of the classification, the class predicted for x . Then, the server and the client privately compute inputs to this polynomial based on x and the thresholds w_i . Finally, the server evaluates the polynomial P privately.

7.1 Polynomial form of a decision tree

Consider that each node of the tree has a boolean variable associated to it. The value of the boolean at a node is 1 if, on input x , one should follow the right branch, and 0 otherwise. For example, denote the boolean variable at the root of the tree by b_1 . The value of b_1 is 1 if $x_1 \leq w_1$ (recall Figure 2), and 0 otherwise.

We construct a polynomial P that, on input all these boolean variables and the value of each class at a leaf node, outputs the class predicted for x . The idea is that P is a sum of terms, where each term (say t) corresponds to a path in the tree from root to a leaf node (say c). A term t evaluates to c iff x is classified along that path in T , else it evaluates to zero. Hence, the term corresponding to a path in the tree is naturally the multiplication of the boolean variables on that path and the class at the leaf node. For example, for the tree in Figure 3, P is $P(b_1, b_2, b_3, b_4, c_1, \dots, c_5) = b_1(b_3 \cdot (b_4 \cdot c_5 + (1 - b_4) \cdot c_4) + (1 - b_3) \cdot c_3) + (1 - b_1)(b_2 \cdot c_2 + (1 - b_2) \cdot c_1)$.

We now present \mathcal{F} , a recursive procedure for constructing P given a binary decision tree T :





Figure 3: Decision tree with booleans



7.2 Private evaluation of a polynomial

Let us first explain how to compute the values of the boolean variables securely. Let n be the number of nodes in the tree and n_{leaves} be the number of leaves in the tree. These values must remain unknown to the server because they leak information about x . For each boolean variable b_i , the server and the client engage in the comparison protocol to compare w_i and the corresponding attribute of x . As a result, the server obtains $[b_i]$ for $i \in 1 \dots n$; the server then changes the encryption of these values to FHE using Protocol 2, thus obtaining $\llbracket b_i \rrbracket$.

The server evaluates P on $(\llbracket b_1 \rrbracket, \dots, \llbracket b_n \rrbracket)$ using the homomorphic properties of FHE. In most cases, FHE evaluation is very slow, but we succeed to make it efficient through a combination of techniques we now discuss. To understand these techniques, recall that a typical FHE evaluation happens over a circuit whose gates are modular addition and multiplication. The performance of FHE depends a lot on the depth of multiplications in this circuit.

First, we use a *leveled* FHE scheme: a scheme that supports only an a priori fixed multiplicative depth instead of an arbitrary such depth. As long as this depth is small, such a scheme is much faster than a *full* FHE scheme.

Second, we ensure that the multiplicative depth is very small using a tree-based evaluation. If h_{max} is the maximum height of the decision tree, then P has a term $a_1 \cdot \dots \cdot a_{h_{\text{max}}}$. If we evaluate this term naively with FHE, we multiply these values sequentially. This yields a multiplicative depth of h_{max} , which makes FHE slow for common h_{max} values. Instead, we construct a binary tree over these values and multiply them in pairs based on the structure of this tree. This results in a multiplicative depth of $\log_2 h_{\text{max}}$ (e.g., 4), which makes FHE evaluation significantly more efficient.

Finally, we use \mathbb{F}_2 as the plaintext space and SIMD slots for parallelism. FHE schemes are significantly faster when the values encrypted are bits (namely, in \mathbb{F}_2); however, P contains classes (e.g., c_1) which are usually more than a bit in length. To enable computing P over F_2 , we represent each class in binary. Let $l = \lceil \log_2 k \rceil$ (k is the number of classes) be the number of bits needed to represent a class. We evaluate P l times, once for each of the l bits of a class. Concretely, the j -th evaluation of P takes as input b_1, \dots, b_n and for each leaf node c_i , its j -th bit c_{ij} . The result is $P(b_1, \dots, b_n, c_{1j}, c_{2j}, \dots, c_{n_{\text{leaves}j}})$, which represents the j -th bit of the outcome class. Hence, we need to run the FHE

evaluation l times.

To avoid this factor of l , the idea is to use a nice feature of FHE called *SIMD slots* (as described in [SV11]): these allow encrypting multiple bits in a single ciphertext such that any operation applied to the ciphertext gets applied in parallel to each of the bits. Hence, for each class c_j , the server creates an FHE ciphertext $\llbracket c_{j0}, \dots, c_{jl-1} \rrbracket$. For each node b_i , it creates an FHE ciphertext $\llbracket b_i, \dots, b_i \rrbracket$ by simply repeating the b_i value in each slot. Then, the server runs one FHE evaluation of P over all these ciphertexts and obtains $\llbracket c_{o0}, \dots, c_{ol-1} \rrbracket$ where c_o is the outcome class. Hence, instead of l FHE evaluations, the server runs the evaluation only once. This results in a performance improvement of $\log k$, a factor of 2 and more in our experiments. We were able to apply SIMD slots parallelism due to the fortunate fact that the same polynomial P had to be computed for each slot.

Finally, evaluating the decision tree is done using $2n$ FHE multiplications and $2n$ FHE additions where n is the number of criteria. The evaluation circuit has multiplication depth $\lceil \log_2(n) + 1 \rceil$.

7.3 Formal description

Protocol 6 describes the resulting protocol.

Protocol 6 Decision Tree Classifier

Client's (C) Input: $x = (x_1, \dots, x_n) \in \mathbb{Z}^n$, secret keys $\text{SK}_{QR}, \text{SK}_{FHE}$

Server's (S) Input: The public keys $\text{PK}_{QR}, \text{PK}_{FHE}$, the model as a decision tree, including the n thresholds $\{w_i\}_{i=1}^n$.

Client's Output: The value of the leaf of the decision tree associated with the inputs b_1, \dots, b_n .

- 1: S produces an n -variate polynomial P as described in section 7.1.
 - 2: S and C interact in the comparison protocol, so that S obtains $[b_i]$ for $i \in [1 \dots n]$ by comparing w_i to the corresponding attribute of x .
 - 3: Using Protocol 2, S changes the encryption from QR to FHE and obtains $\llbracket b_1 \rrbracket, \dots, \llbracket b_n \rrbracket$.
 - 4: To evaluate P , S encrypts the bits of each category c_i using FHE and SIMD slots, obtaining $\llbracket c_{i1}, \dots, c_{il} \rrbracket$. S uses SIMD slots to compute homomorphically $\llbracket P(b_1, \dots, b_n, c_{10}, \dots, c_{n\text{leaves}0}), \dots, P(b_1, \dots, b_n, c_{1l-1}, \dots, c_{n\text{leaves}l-1}) \rrbracket$. It rerandomizes the resulting ciphertext using FHE's rerandomization function, and sends the result to the client.
 - 5: C decrypts the result as the bit vector (v_0, \dots, v_{l-1}) and outputs $\sum_{i=0}^{l-1} v_i \cdot 2^i$.
-

Proposition 7.1. *Protocol 6 is secure in the honest-but-curious model.*

Proof intuition. The proof is in Appendix C, but we give some intuition here. During the comparison protocol, the server only learns encrypted bits, so it learns nothing about x . During FHE evaluation, it similarly learns nothing about the input due to the security of FHE. The client does not learn the structure of the tree because the server performs the evaluation of the polynomial. Similarly, the client does not learn the bits at the nodes in the tree because of the security of the comparison protocol.

The interactions between the client and the server are due to the comparisons almost exclusively: the decision tree evaluation does not need any interaction but sending the encrypted result of the evaluation.

8 Combining classifiers with AdaBoost

AdaBoost is a technique introduced in [FS97]. The idea is to combine a set of *weak* classifiers $h_i(x) : \mathbb{R}^d \mapsto \{-1, +1\}$ to obtain a better classifier. The AdaBoost algorithm chooses t scalars $\{\alpha_i\}_{i=1}^t$ and constructs a strong classifier as:

$$H(x) = \text{sign} \left(\sum_{i=1}^t \alpha_i h_i(x) \right)$$

If each of the $h_i(\cdot)$'s is an instance of a classifier supported by our protocols, then given the scalars α_i , we can easily and securely evaluate $H(x)$ by simply composing our building blocks. First, we run the secure protocols for each of

Bit size	A Computation	B Computation	Total Time	Communication	Interactions
10	5.03 ms	12.38 ms	98.4 ms	5.26 kB	3
20	8.34 ms	17.3 ms	107 ms	10.4 kB	3
32	13.70 ms	17.58 ms	112 ms	16.5 kB	3
64	26.15 ms	39.03 ms	149 ms	32.9 kB	3

Table 3: Comparison with unencrypted input protocols evaluation.

Protocol	Bit size	Computation		Total Time	Communication	Interactions
		Party A	Party B			
Comparison	64	38.07 ms	20.99 ms	258.5 ms	33.41 kB	6
Reversed Comp.	64	25.83 ms	36.93 ms	292.1 ms	33.41 kB	6

Table 4: Comparison with encrypted input protocols evaluation.

h_i , except that the server keeps the intermediate result, the outcome of $h_i(x)$, encrypted using one of our comparison protocols (Rows 2 or 4 of Table 2). Second, if necessary, we convert them to Paillier’s encryption scheme with Protocol 2, and combine these intermediate results using Paillier’s additive homomorphic property as in the dot product protocol Protocol 3. Finally, we run the comparison over encrypted data algorithm to compare the result so far with zero, so that the client gets the final result.

Party A Computation	Party B Computation	Total Time	Communication	Interactions
47.0 ms	232 ms	549 ms	481.2 kB	2

Table 5: Change encryption scheme protocol evaluation.

9 Implementation

We have implemented the protocols and the classifiers in C++ using GMP⁴, Boost, Google’s Protocol Buffers⁵, and HELib [Hal13] for the FHE implementation.

The code is written in a modular way: all the elementary protocols defined in Section 4 can be used as black boxes with minimal developer effort. Thus, writing secure classifiers comes down to invoking the right API calls to the protocols. For example, for the linear classifier, the client simply calls a key exchange protocol to setup the various keys, followed by the dot product protocol, and then the comparison of encrypted data protocol to output the result, as shown in Figure 4.

10 Evaluation

To evaluate our work, we answer the following questions: (i) can our building blocks be used to construct other classifiers in a modular way (Section 10.1), (ii) what is the performance overhead of our building blocks (Section 10.3), and (iii) what is the performance overhead of our classifiers (Section 10.4)?

10.1 Using our building blocks library

Here we demonstrate that our building blocks library can be used to build other classifiers modularly and that it is a useful contribution by itself. We will construct a multiplexer and a face detector. A face detection algorithm over encrypted data already exists [AB06, AB07], so our construction here is not the first such construction, but it serves as a proof of functionality for our library.

⁴<http://gmplib.org/>

⁵<https://code.google.com/p/protobuf/>


```

bool Linear_Classifier_Client::run()
{
    exchange_keys();

    // values_ is a vector of integers
    // compute the dot product
    mpz_class v = compute_dot_product(values_);
    mpz_class w = 1; // encryption of 0

    // compare the dot product with 0
    return enc_comparison(v, w, bit_size_, false);
}

void Linear_Classifier_Server_session::
    run_session()
{
    exchange_keys();

    // enc_model_ is the encrypted model vector
    // compute the dot product
    help_compute_dot_product(enc_model_, true);

    // help the client to get
    // the sign of the dot product
    help_enc_comparison(bit_size_, false);
}

```

Figure 4: Implementation example: a linear classifier

10.1.1 Building a multiplexer classifier

A multiplexer is the following generalized comparison function:

$$f_{\alpha,\beta}(a, b) = \begin{cases} \alpha & \text{if } a > b \\ \beta & \text{otherwise} \end{cases}$$

We can express $f_{\alpha,\beta}$ as a linear combination of the bit $d = (a \leq b)$:

$$f_{\alpha,\beta}(d) = d \cdot \beta + (1 - d) \cdot \alpha = \alpha + d \cdot (\beta - \alpha).$$

To implement this classifier privately, we compute $\llbracket d \rrbracket$ by comparing a and b , keeping the result encrypted with QR, and then changing the encryption scheme (cf. Section 4.3) to Paillier.

Then, using Paillier’s homomorphism and knowledge of α and β , we can compute an encryption of $f_{\alpha,\beta}(d)$:

$$\llbracket f_{\alpha,\beta}(d) \rrbracket = \llbracket \alpha \rrbracket \cdot \llbracket d \rrbracket^{\beta - \alpha}.$$

10.1.2 Viola and Jones face detection

The Viola and Jones face detection algorithm [VJ01] is a particular case of an AdaBoost classifier. Denote by X an image represented as an integer vector and x a particular detection window (a subset of X ’s coefficients). The *strong* classifier H for this particular detection window is

$$H(x) = \text{sign} \left(\sum_{i=1}^t \alpha_i h_i(x) \right)$$

where the h_t are weak classifiers of the form $h_i(x) = \text{sign}(\langle x, y_i \rangle - \theta_i)$.

In our setting, Alice owns the image and Bob the classifier (e.g. the vectors $\{y_i\}$ and the scalars $\{\theta_i\}$ and $\{\alpha_i\}$). Neither of them wants to disclose their input to the other party. Thanks to our building blocks, Alice can run Bob’s classifier on her image without her learning anything about the parameters and Bob learning any information about her image.

The weak classifiers can be seen as multiplexers; with the above notation, we have $h_t(x) = f_{1,-1}(\langle x, y_t \rangle - \theta_t)$.

Using the elements of Section 10.1.1, we can easily compute the encrypted evaluation of every one of these weak classifiers under Paillier, and then, as described in Section 8, compute the encryption of $H(x)$.

10.2 Performance evaluation setup

Our performance evaluations were run using two desktop computers each with identical configuration: two Intel i7 (64 bit) processors for a total 8 cores running at 3.4 GHz and 8 GB RAM. Since the machines were on the same network,

we inflated the roundtrip time for a packet to be 40 ms to mimic real network latency. We used 1024-bit cryptographic keys, and chose the statistical security parameter λ to be 100. When using HELib, we use 80 bits of security, which corresponds to a 1024-bit asymmetric key.

10.3 Building blocks performance

We examine performance in terms of computation time at the client and server, communication bandwidth, and also number of interactions (round trips). We can see that all these protocols are efficient, with a runtime on the order of milliseconds.

10.3.1 Comparison protocols

Comparison with unencrypted input. Table 3 gives the running time of the DGK comparison protocol with unencrypted input for various input size. The DGK protocol runs in parallel using four threads for each party.

Comparison with encrypted input. Table 4 presents the performance of the comparison with encrypted inputs protocols, with DGK as underlying comparison protocols.

10.3.2 argmax

Figure 5 presents the running times and the communication overhead of the argmax of encrypted data protocol (*cf.* Section 4.2). The input integers were 64 bit integers.

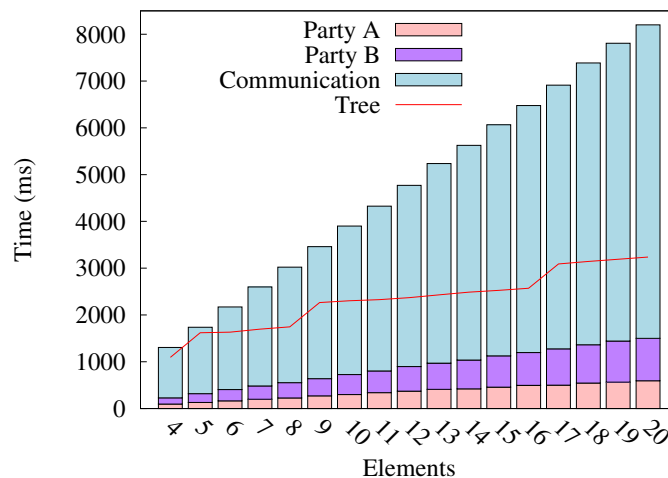


Figure 5: Argmax of encrypted data protocol evaluation. The bars represent the execution of the protocol when the comparisons are executed one after each other, linearly. The line represents the execution when comparisons are executed in parallel, tree-wise.

10.4 Classifier performance

Here we evaluate each of the classifiers described in Sections 5–7. The models are trained non-privately using `scikit-learn`⁶. We used the following datasets from the UCI machine learning repository [BL13]:

1. the Wisconsin Diagnostic Breast Cancer data set,
2. the Wisconsin Breast Cancer (Original) data set, a simplified version of the previous dataset,

⁶<http://scikit-learn.org>

Data set	Model size	Computation		Time per protocol		Total running time	Comm.	Interactions
		Client	Server	Compare	Dot product			
Breast cancer (1)	30	52.0 ms	35.7 ms	281 ms	45.4 ms	326 ms	41.35 kB	7
Credit (3)	47	55.8 ms	44.0 ms	275 ms	48.2 ms	323 ms	45.70 kB	7

(a) Linear Classifier. Time per protocol includes communication.

Data set	Specs.		Computation		Time per protocol		Total running time	Comm.	Interactions
	C	F	Client	Server	Prob. Comp.	Argmax			
Breast Cancer (2)	2	9	122 ms	148 ms	82.1 ms	469 ms	551 ms	77.26 kB	14
Nursery (5)	5	9	434 ms	393 ms	82.9 ms	1563 ms	1646 ms	171.1 kB	42
Audiology (4)	24	70	1557 ms	1931 ms	636 ms	3388 ms	4024 ms	2067 kB	166

(b) Naïve Bayes Classifier. C is the number of classes and F is the number of features. The Prob. Comp. column corresponds to the computation of the probabilities $p(c_i|x)$ (cf. Section 6). Time per protocol includes communication.

Data set	Tree Specs.		Computation		Time per protocol		FHE		Comm.	Interactions
	N	D	Client	Server	Compare	ES Change	Eval.	Decrypt		
Nursery (5)	4	4	991 ms	1756 ms	1474 ms	1709 ms	182 ms	1443 ms	3210 kB	30
ECG (6)	6	4	1485 ms	2595 ms	2309 ms	2627 ms	689 ms	2005 ms	4272 kB	44

(c) Decision Tree Classifier. ES change indicates the time to run the protocol for changing encryption schemes. N is the number of nodes of the tree and D is its depth. Time per protocol includes communication.

Table 6: Classifiers evaluation.

3. Credit Approval data set,
4. Audiology (Standardized) data set,
5. Nursery data set, and
6. ECG (electrocardiogram) classification data from Barni *et al.* [BFK⁺09]

These data sets are scenarios when we want to ensure privacy of the server’s model and client’s input.

Based on the suitability of each classifier, we used data sets 2 and 3 to test the hyperplane decision classifier, sets 1, 4 and 5 for the Naïve Bayes classifier, and sets 5 and 6 for the decision tree classifier.

Table 6 shows the performance results. Our classifiers run in at most a few seconds, which we believe to be practical for sensitive applications. Note that even if the datasets become very large, the size of the model stays the same – the dataset size only affects the training phase which happens on unencrypted data before one uses our classifiers. Hence, the cost of our classification will be the same even for very large data sets.

For the decision tree classifier, we compared our construction to Barni *et al.* [BFK⁺09] on the ECG dataset (by turning their branching program into a decision tree). Their performance is 1765 ms⁷ for the client and 4235 ms for the server with communication cost of 112.2KB. Even though their evaluation does not consider the communication delays, we are still twice as fast for the server and faster for the client. Moreover, we must not forget that we keep the tree private while Barni *et al.* reveals the computation circuit (by revealing the garbled circuit).

10.5 Comparison to generic two-party tools

A set of generic secure two- or multi-party computation tools have been developed, such as TASTY [HKoS⁺10] and Fairplay [MNPS04, BDNP08]. These support general functions, which include our classifiers.

However, they are prohibitively slow for our specific setting. Our efficiency comes from specializing to classification functionality. To demonstrate their performance, we attempted to evaluate the Naïve Bayes classifier with these. We

⁷In Barni *et al.* [BFK⁺09], the evaluation was run over two 3GHz computers directly connected via Gigabit Ethernet. We scaled the given results by $\frac{3}{3.4}$ to get a better comparison basis.

used FairplayMP to generate the circuit for this classifier and then TASTY to run the private computation on the circuit thus obtained. We tried to run the smallest Naïve Bayes instance, the Nursery dataset from our evaluation, which has only 3 possible values for each feature, but we ran out of memory during the circuit generation phase on a powerful machine with 256GB of RAM.

Hence, we had to reduce the classification problem to only 3 classes (versus 5). Then, the circuit generation took more than 2 hours with FairplayMP, and the time to run the classification with TASTY was 413196 msec (with no network delay), which is ≈ 500 times slower than our performance (on the non-reduced classification problem with 5 classes). Thus, our specialized protocols improve performance by orders of magnitude.

11 Conclusion

In this paper, we constructed three major privacy-preserving classifiers as well as provided a library of building blocks that enables constructing other classifiers. We demonstrated the efficiency of our classifiers on real datasets.

References

- [AB06] Shai Avidan and Moshe Butman. Blind vision. In *Computer Vision—ECCV 2006*, pages 1–13. Springer, 2006.
- [AB07] Shai Avidan and Moshe Butman. Efficient methods for privacy preserving face detection. In *Advances in Neural Information Processing Systems: Proceedings of the 2006 Conference*, page 57, 2007.
- [AD01] Mikhail J Atallah and Wenliang Du. Secure multi-party computational geometry. In *Algorithms and Data Structures*, pages 165–179. Springer, 2001.
- [ALSZ13] Gilad Asharov, Yehuda Lindell, Thomas Schneider, and Michael Zohner. More efficient oblivious transfer and extensions for faster secure computation. In *CCS*, pages 535–548, 2013.
- [BDMN05] Avrim Blum, Cynthia Dwork, Frank McSherry, and Kobbi Nissim. Practical privacy: the sulq framework. In *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 128–138. ACM, 2005.
- [BDNP08] Assaf Ben-David, Noam Nisan, and Benny Pinkas. Fairplaymp: A system for secure multi-party computation. In *CCS*, pages 17–21, 2008.
- [BFK⁺09] Mauro Barni, Pierluigi Failla, Vladimir Kolesnikov, Riccardo Lazzeretti, Ahmad-Reza Sadeghi, and Thomas Schneider. Secure evaluation of private linear branching programs with medical applications. In *Computer Security—ESORICS 2009*, pages 424–439. Springer, 2009.
- [BFL⁺09] Mauro Barni, Pierluigi Failla, Riccardo Lazzeretti, Annika Paus, A-R Sadeghi, Thomas Schneider, and Vladimir Kolesnikov. Efficient privacy-preserving classification of ecg signals. In *Information Forensics and Security, 2009. WIFS 2009. First IEEE International Workshop on*, pages 91–95. IEEE, 2009.
- [BGV12] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325, 2012.
- [BHKR13] Mihir Bellare, Viet Tung Hoang, Sriram Keelveedhi, and Phillip Rogaway. Efficient garbling from a fixed-key blockcipher. In *IEEE SP*, pages 478–492, 2013.
- [BL13] K. Bache and M. Lichman. UCI machine learning repository, 2013.
- [BLN13] Joppe W. Bos, Kristin Lauter, and Michael Naehrig. Private predictive analysis on encrypted medical data, 2013.
- [BN06] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 1. springer New York, 2006.
- [Can98] Ran Canetti. Security and composition of multi-party cryptographic protocols. *JOURNAL OF CRYPTOLOGY*, 13:2000, 1998.
- [CMS11] Kamalika Chaudhuri, Claire Monteleoni, and Anand D. Sarwate. Differentially private empirical risk minimization. *J. Mach. Learn. Res.*, 12, 2011.
- [DGK07] Ivan Damgård, Martin Geisler, and Mikkel Krøigaard. Efficient and secure comparison for on-line auctions. In *Information Security and Privacy*, pages 416–430. Springer, 2007.
- [DGK09] Ivan Damgård, Martin Geisler, and Mikkel Kroigard. A correction to ‘efficient and secure comparison for on-line auctions’. *International Journal of Applied Cryptography*, 1(4):323–324, 2009.
- [DHC04] Wenliang Du, Yungshiang S Han, and Shigang Chen. Privacy-preserving multivariate statistical analysis: Linear regression and classification. In *Proceedings of the 4th SIAM International Conference on Data Mining*, volume 233. Lake Buena Vista, Florida, 2004.

- [EFG⁺09] Zekeriya Erkin, Martin Franz, Jorge Guajardo, Stefan Katzenbeisser, Inald Lagendijk, and Tomas Toft. Privacy-preserving face recognition. In *Privacy Enhancing Technologies*, pages 235–253. Springer, 2009.
- [FS97] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- [Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178, 2009.
- [GLLM05] Bart Goethals, Sven Laur, Helger Lipmaa, and Taneli Mielikäinen. On private scalar product computation for privacy-preserving data mining. In *Information Security and Cryptology–ICISC 2004*, pages 104–120. Springer, 2005.
- [GLN13] Thore Graepel, Kristin Lauter, and Michael Naehrig. ML confidential: Machine learning on encrypted data. In *Information Security and Cryptology–ICISC 2012*, pages 1–21. Springer, 2013.
- [GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *STOC*, pages 365–377. ACM, 1982.
- [GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game. In *STOC*, pages 218–229, 1987.
- [Gol04] Oded Goldreich. *Foundations of Cryptography - Basic Applications*. Cambridge University Press, 2004.
- [Hal13] Shai Halevi. Helib - an implementation of homomorphic encryption. <https://github.com/shaih/HElib>, 2013.
- [HKoS⁺10] Wilko Henecka, Stefan Kögl, Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. Tasty: Tool for automating secure two-party computations. In *CCS*, pages 451–462, 2010.
- [IPS08] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Founding cryptography on oblivious transfer – efficiently. In David Wagner, editor, *Advances in Cryptology – CRYPTO 2008*, volume 5157 of *Lecture Notes in Computer Science*, pages 572–591. Springer Berlin Heidelberg, 2008.
- [Kil05] Eike Kiltz. Unconditionally secure constant round multi-party computation for equality, comparison, bits and exponentiation. *IACR Cryptology ePrint Archive*, 2005:66, 2005.
- [KSS] Vladimir Kolesnikov, Ahmad-Reza Sadeghi, and Thomas Schneider. A systematic approach to practically efficient general two-party secure function evaluation protocols and their modular design. page 2013.
- [KSS09] Vladimir Kolesnikov, Ahmad-Reza Sadeghi, and Thomas Schneider. Improved garbled circuit building blocks and applications to auctions and computing minima. In *In Cryptology and Network Security (CANS)*, 2009.
- [LLM06] Sven Laur, Helger Lipmaa, and Taneli Mielikäinen. Cryptographically private support vector machines. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 618–624. ACM, 2006.
- [LP00] Yehuda Lindell and Benny Pinkas. Privacy preserving data mining. In *Advances in Cryptology—CRYPTO 2000*, pages 36–54. Springer, 2000.
- [LP07] Yehuda Lindell and Benny Pinkas. An efficient protocol for secure two-party computation in the presence of malicious adversaries. In Moni Naor, editor, *Advances in Cryptology - EUROCRYPT 2007*, volume 4515 of *Lecture Notes in Computer Science*, pages 52–78. Springer Berlin Heidelberg, 2007.
- [LP08] Yehuda Lindell and Benny Pinkas. Secure multiparty computation for privacy-preserving data mining, 2008.

- [LP09] Yehuda Lindell and Benny Pinkas. A proof of security of Yao’s protocol for two-party computation. *J. Cryptol.*, 22:161–188, April 2009.
- [LT05] Hsiao-Ying Lin and Wen-Guey Tzeng. An efficient solution to the millionaires’ problem based on homomorphic encryption. In *Applied Cryptography and Network Security*, pages 456–466. Springer, 2005.
- [MNPS04] Dahlia Malkhi, Noam Nisan, Benny Pinkas, and Yaron Sella. Fairplay-secure two-party computation system. In *USENIX Security Symposium*, pages 287–302, 2004.
- [Pai99] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *EUROCRYPT*, pages 223–238, 1999.
- [RR07] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *NIPS*, 2007.
- [SG11] Anima Singh and John Gutttag. Cardiovascular risk stratification using non-symmetric entropy-based classification trees. In *NIPS workshop on personalized medicine*, 2011.
- [SG13] Anima Singh and John Gutttag. Leveraging hierarchical structure in diagnostic codes for predicting incident heart failure. In *ICML workshop on role of machine learning in transforming healthcare*, 2013.
- [SSW10] Ahmad-Reza Sadeghi, Thomas Schneider, and Immo Wehrenberg. Efficient privacy-preserving face recognition. In *Information, Security and Cryptology–ICISC 2009*, pages 229–244. Springer, 2010.
- [SV11] Nigel P. Smart and Frederik Vercauteren. Fully homomorphic SIMD operations. Cryptology ePrint Archive, Report 2011/133, 2011.
- [TRMP12] Sebastian Tschiatschek, Peter Reinprecht, Manfred Mücke, and Franz Pernkopf. Bayesian network classifiers with reduced precision parameters. In *Machine Learning and Knowledge Discovery in Databases*, pages 74–89. Springer, 2012.
- [Veu11] Thijs Veugen. Comparing encrypted data. <http://msp.ewi.tudelft.nl/sites/default/files/Comparing%20encrypted%20data.pdf>, 2011.
- [VJ01] Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, volume 1, pages I–511. IEEE, 2001.
- [VKC08] Jaideep Vaidya, Murat Kantarcioğlu, and Chris Clifton. Privacy-preserving naive bayes classification. *The International Journal on Very Large Data Bases*, 17(4):879–898, 2008.
- [WGH12] Jenna Wiens, John Gutttag, and Eric Horvitz. Learning evolving patient risk processes for c. diff colonization. In *ICML*, 2012.
- [WY04] Rebecca Wright and Zhiqiang Yang. Privacy-preserving bayesian network structure computation on distributed heterogeneous data. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 713–718. ACM, 2004.
- [Yao82] Andrew C. Yao. Protocols for secure computations. In *FOCS*, pages 160–164, 1982.
- [ZW05] Zhiqiang Yang, Sheng Zhong and Rebecca N Wright. Privacy-preserving classification of customer data without loss of accuracy. In *SIAM International Conference on Data Mining (SDM), Newport Beach*. Citeseer, 2005.

A Comparison protocols

A.1 Comparison with unencrypted inputs

Our protocol for comparing with encrypted inputs is Protocol 7 and here is some intuition. We follow the main idea from Veugen [Ve11] (found in Section 2.1): compute $2^l + b - a$ (over encrypted data) and check the $l + 1$ -th bit (the bit corresponding to the power 2^l). If it is 1, it means that $b \geq a$, else $b < a$.

We also assume that the encryption scheme is additively homomorphic. In [Ve11] (Section 2.1), Veugen presents a solution for a similar problem except that A only gets the encrypted bit, not in the clear. So we modify his protocol in Protocol 7.

In the description of protocol 7, N is the modulus associated with Paillier's cryptosystem.

Protocol 7 Comparing encrypted data

Input A: $\llbracket a \rrbracket$ and $\llbracket b \rrbracket$, the bit length l of a and b , the secret key SK_{QR} , public key PK_P

Input B: Secret key SK_P , public key PK_{QR} , the bit length l

Output A: $(a \leq b)$

- 1: A: $\llbracket x \rrbracket \leftarrow \llbracket b \rrbracket \cdot \llbracket 2^l \rrbracket \cdot \llbracket a \rrbracket^{-1} \bmod N^2$
 - 2: A chooses a random number $r \leftarrow (0, 2^{\lambda+l}) \cap \mathbb{Z}$
 - 3: A: $\llbracket z \rrbracket \leftarrow \llbracket x \rrbracket \cdot \llbracket r \rrbracket \bmod N^2$ ▷ Blind x
 - 4: A sends $\llbracket z \rrbracket$ to B
 - 5: B decrypts $\llbracket z \rrbracket$
 - 6: A: $c \leftarrow r \bmod 2^l$
 - 7: B: $d \leftarrow z \bmod 2^l$
 - 8: With A, B privately computes the encrypted bit $\llbracket t' \rrbracket$ such that $t = (d < c)$ using DGK
 - 9: A encrypts r_l and sends $\llbracket r_l \rrbracket$ to B
 - 10: B encrypts z_l
 - 11: B: $\llbracket t \rrbracket \leftarrow \llbracket t' \rrbracket \cdot \llbracket z_l \rrbracket \cdot \llbracket r_l \rrbracket$
 - 12: B: sends $\llbracket t \rrbracket$ to A
 - 13: A decrypts and outputs t
-

We will show the correctness of the protocol and then give a proof of security in the honest-but-curious model using modular composition. For the correctness, we just modify the proof of [Ve11].

Proposition A.1. *Protocol 7 is correct and secure in the honest-but-curious model.*

See proof in Appendix C.

A.2 Reversed encrypted comparison

We constructed Protocol 8 which is the same as Protocol 7, except that the roles of A and B are exchanged in Steps 8–13.

Proposition A.2. *Protocol 8 is secure in the honest-but-curious model.*

The proof is in Appendix C.

Protocol 8 Reversed comparing encrypted data

Input A: $\llbracket a \rrbracket$ and $\llbracket b \rrbracket$, public keys PK_{QR} and PK_P

Input B: Secret keys SK_P and SK_{QR}

Output B: $(a \leq b)$

Run Steps 1–7 of Protocol 7.

8: With B, A privately computes the encrypted bit $[t']$ such that $t' = (d < c)$ using DGK

9: B encrypts z_l and sends $[z_l]$ to A

10: A encrypts r_l

11: A: $[t] \leftarrow [t'] \cdot [z_l] \cdot [r_l]$

12: A: sends $[t]$ to B

13: B decrypts and outputs t

B Preliminaries for proofs

B.1 Secure two-party computation framework

All our protocols are two-party protocols, which we label as party A and party B. In order to show that they do private computations, we work in the honest-but-curious (semi-honest) model as described in [Gol04].

Let $f = (f_A, f_B)$ be a (probabilistic) polynomial function and Π a protocol computing f . A and B want to compute $f(a, b)$ where a is A's input and b is B's input, using Π and with the security parameter λ . The view of party A during the execution of Π is the tuple $V_A(\lambda, a, b) = (1^\lambda; a; r^A; m_1^A, \dots, m_t^A)$ where r is A's random tape and m_1^A, \dots, m_t^A are the messages received by A. We define the view of B similarly. We also define the outputs of parties A and B for the execution of Π on input (a, b) as $\text{Output}_A^\Pi(\lambda, a, b)$ and $\text{Output}_B^\Pi(\lambda, a, b)$, and the global output as $\text{Output}^\Pi(\lambda, a, b) = (\text{Output}_A^\Pi(\lambda, a, b), \text{Output}_B^\Pi(\lambda, a, b))$.

To ensure security, we have to show that whatever A can compute from its interactions with B can be computed from its input and output, which leads us to the following security definition.

Definition B.1. *The two-party protocol Π securely computes the function f if there exists two probabilistic polynomial time algorithms S_A and S_B such that for every possible input a, b of f ,*

$$\begin{aligned} \{S_A(1^\lambda, a, f_A(a, b)), f(a, b)\} &\equiv_c \\ &\{V_A(\lambda, a, b), \text{Output}^\Pi(\lambda, a, b)\} \end{aligned}$$

and

$$\begin{aligned} \{S_B(1^\lambda, a, f_B(a, b)), f(a, b)\} &\equiv_c \\ &\{V_B(\lambda, a, b), \text{Output}^\Pi(\lambda, a, b)\} \end{aligned}$$

where \equiv_c means computational indistinguishability against probabilistic polynomial time adversaries with negligible advantage in the security parameter λ .

To simplify the notation (and the proofs), hereinafter we omit the security parameter. As we mostly consider deterministic functions f , we can simplify the distributions we want to show being indistinguishable (see [Gol04]): when f is deterministic, to prove the security of Π that computes f , we only have to show that

$$\begin{aligned} S_A(a, f_A(a, b)) &\equiv_c V_A(a, b) \\ S_B(b, f_B(a, b)) &\equiv_c V_B(a, b) \end{aligned}$$

Unless written explicitly, we will always prove security using this simplified definition.

B.2 Modular Sequential Composition

In order to ease the proofs of security, we use sequential modular composition, as defined in [Can98]. The idea is that the parties run a protocol Π and use calls to an ideal functionality f in Π (e.g. A and B compute f privately by sending their inputs to a trusted third party and receiving the result). If we can show that Π respects privacy in the honest-but-curious model and if we have a protocol ρ that privately computes f in the same model, then we can replace the ideal calls for f by the execution of ρ in Π ; the new protocol, denoted Π^ρ is then secure in the honest-but-curious model.

We call *hybrid model with ideal access to f_1, \dots, f_m* or *(f_1, \dots, f_m) -hybrid model* the semi-honest model augmented with an incorruptible trusted party T for evaluating functionalities f_1, \dots, f_m . The parties run a protocol Π that contain calls to T for the evaluation of one of f_1, \dots, f_m . For each call, each party sends its input and wait until the trusted party sends the output back. We emphasize on the fact that the parties must not communicate until receiving T 's output (we consider only *sequential* composition). Ideal calls to the trusted party can be done several times, even for the same function, but each call is independent: T does not maintain state between two calls.

Let Π be a two-party protocol in the (f_1, \dots, f_m) -hybrid model. Let ρ_1, \dots, ρ_m be real protocols (i.e. protocols in the semi-honest model) computing f_1, \dots, f_m and define $\Pi^{\rho_1, \dots, \rho_m}$ as follows. All ideal calls of Π to the trusted party for f_i is replaced by a real execution of ρ_i : if party P_j has to compute f_i with input x_j , P_j halts, starts an execution of ρ_i with the other parties, gets the result β_j when ρ_i concludes, and continues as if β_j was received from T .

Theorem B.2. [Can98] (Theorem 5) restated as in [LP08] (Theorem 3) – *Let f_1, \dots, f_m be two-party probabilistic polynomial time functionalities and ρ_1, \dots, ρ_m protocols that compute respectively f_1, \dots, f_m in the presence of semi-honest adversaries.*

Let g be a two-party probabilistic polynomial time functionality and Π a protocol that securely computes g in the (f_1, \dots, f_m) -hybrid model in the presence of semi-honest adversaries.

Then $\Pi^{\rho_1, \dots, \rho_m}$ securely computes g in the presence of semi-honest adversaries.

B.3 Cryptographic assumptions

Assumption 1. (Quadratic Residuosity Assumption – from [GM82]) *Let $N = p \times q$ be the product of two distinct odd primes p and q . Let \mathbb{QR}_N be the set of quadratic residues modulo N and \mathbb{QNR}_N be the set of quadratic non residues (i.e. $x \in \mathbb{QNR}_N$ if x is not a square modulo N and its Jacobi symbol is 1).*

$\{(N, \mathbb{QR}_N) : |N| = \lambda\}$ and $\{(N, \mathbb{QNR}_N) : |N| = \lambda\}$ are computationally indistinguishable with respect to probabilistic polynomial time algorithms.

Assumption 2. (Decisional Composite Residuosity Assumption – from [Pai99]) *Let $N = p \times q, |N| = \lambda$ be the product of two distinct odd primes p and q . A number z is said to be a N -th residue modulo N^2 if there exists a number $y \in \mathbb{Z}_{N^2}$*

$$z = y^N \pmod{N^2}$$

N -th residues are computationally indistinguishable from non N -th residues with respect to probabilistic polynomial time algorithms.

For further explanations about the last assumption, used for the FHE scheme, we refer the reader to [BGV12].

Assumption 3. (RLWE) *For security parameter λ , let $f(x) = x^d + 1$ where d is a power of 2. Let $q \geq 2$ be an integer. Let $R = \mathbb{Z}[x]/(f(x))$ and let $R_q = R/qR$. Let χ be a distribution over R . The $\text{RLWE}_{d,q,\chi}$ problem is to distinguish between two distributions: In the first distribution, one samples (a_i, b_i) uniformly from R_q^2 . In the second distribution, one first draws $s \leftarrow R_q$ uniformly and then samples $(a_i, b_i) \in R_q^2$ by sampling $a_i \leftarrow R_q$ uniformly, $e_i \leftarrow \chi$, and setting $b_i = a_i \cdot s + e_i$.*

The $\text{RLWE}_{d,q,\chi}$ assumption is that the $\text{RLWE}_{d,q,\chi}$ problem is infeasible.

C Proofs

C.1 Comparison protocols

Proof of Proposition A.1. Correctness As a and b are l bits integers, $x = 2^l + b - a$ is a $l + 1$ bits integer and its most significant bit (the $l + 1$ -th bit) is 1 iff $a \leq b$. What protocol 7 actually does is computing this bit. The computations are done over encrypted data, using Paillier's encryption scheme. In the rest of the proof, we will do as if the data were not encrypted under Paillier. The correctness will hold as long as we do not experience carry-overs modulo N . In particular, this implies that $l + 1 + \lambda < \log_2 N$. For operations over bits using QR, we don't have this problem as we are operating on \mathbb{F}_2 .

Again, since x is a $l + 1$ bit number, its most significant bit is $x \div 2^l$ where \div denotes the integer division. We have $x = 2^l(x \div 2^l) + (x \bmod 2^l)$ where $0 \leq (x \bmod 2^l) < 2^l$. As $z = x + r$,

$$\begin{aligned} z &= 2^l(z \div 2^l) + (z \bmod 2^l) \\ &= 2^l((x \div 2^l) + (r \div 2^l)) + ((x \bmod 2^l) + (r \bmod 2^l)) \end{aligned}$$

Hence, $z \div 2^l = x \div 2^l + r \div 2^l$ if $(x \bmod 2^l) + (r \bmod 2^l) < 2^l$ and $z \div 2^l = (x \div 2^l) + (r \div 2^l) + 1$ otherwise. More generally, $z \div 2^l = (x \div 2^l) + (r \div 2^l) + t'$ where $t' = 0 \Leftrightarrow (x \bmod 2^l) + (r \bmod 2^l) < 2^l$.

We can also notice that, if $t' = 0$, $z \bmod 2^l = (x \bmod 2^l) + (r \bmod 2^l)$ and $z \bmod 2^l = (x \bmod 2^l) + (r \bmod 2^l) - 2^l$ otherwise. As a consequence,

$$\begin{aligned} t' = 0 &\Leftrightarrow z \bmod 2^l = (x \bmod 2^l) + (r \bmod 2^l) \\ &\Leftrightarrow z \bmod 2^l \geq (r \bmod 2^l) \end{aligned}$$

In the end, as $x \div 2^l$ is either 0 or 1, we can compute everything modulo 2

$$\begin{aligned} x \div 2^l &= (z \div 2^l) - (r \div 2^l) - t' \pmod{2} \\ &= z_l \oplus r_l \oplus t' \end{aligned}$$

Security We suppose that the encrypted bit $[t']$ is ideally computed (using calls to a trusted party in the hybrid model). We show that the protocol is secure in this model and conclude using the sequential modular composition theorem.

A's view is $V_A = ([a], [b], l, \text{SK}_{QR}, \text{PK}_P; r, \text{coins}; [t])$ where SK_{QR} is the secret key for the QR cryptosystem, PK_P is the public key for Paillier's cryptosystem, and coins are the random coins used for the encryptions of 2^l , r and r_l . Given $([a], [b], l, \text{SK}_{QR}, \text{PK}_P, a \leq b)$, we build the simulator S_A :

1. Compute $[\tilde{t}]$ an encryption of the bit $(a \leq b)$ under QR.
2. Pick $\tilde{r} \leftarrow (0, 2^{\lambda+l}) \cap \mathbb{Z}$.
3. Let $\widetilde{\text{coins}}$ be random coins for two Paillier encryptions and one QR encryption.
4. Output $([a], [b], l, \text{SK}_{QR}, \text{PK}_P; \tilde{r}, \widetilde{\text{coins}}; [\tilde{t}])$

The distributions $V_A([a], [b], l, \text{SK}_{QR}, \text{PK}_P, a \leq b)$ and $S_A([a], [b], \text{SK}_{QR}, \text{PK}_P, a \leq b)$ are exactly the same because the randomness is taken from the same distribution in both cases, and the QR cyphertext encrypts the same bit.

B's view is $V_B = (\text{PK}_{QR}, \text{SK}_P, l, [z]; \text{coins}; [t'], [r_l])$ where coins are the random coins used for the encryption of z_l . The simulator $S_B(\text{PK}_{QR}, \text{SK}_P, l)$ runs as follows:

1. Pick $\tilde{z} \leftarrow (0, 2^{\lambda+l}) \cap \mathbb{Z}$.
2. Encrypt \tilde{z} under Paillier: $[\tilde{z}]$.
3. Generate $[\tilde{t}']$ and $[\tilde{r}_l]$, two encryptions of random bits under QR
4. Let $\widetilde{\text{coins}}$ be random coins for one QR encryption.
5. Output $(\text{PK}_{QR}, \text{SK}_P, l, [\tilde{z}]; \widetilde{\text{coins}}; [\tilde{t}'], [\tilde{r}_l])$

The random tapes coins and $\widetilde{\text{coins}}$ are generated in the exact same manner and independently from any other parameter, so

$$\begin{aligned} & (\text{PK}_{QR}, \text{SK}_P, \llbracket \tilde{z} \rrbracket; \widetilde{\text{coins}}; [\tilde{t}'][\tilde{r}_l]) \\ &= (\text{PK}_{QR}, \text{SK}_P, \llbracket \tilde{z} \rrbracket; \text{coins}; [\tilde{t}'][\tilde{r}_l]) \end{aligned}$$

Recall that $z = x + r \bmod N$ where x is an l bits integer and r is an $l + \lambda$ bits integer. But as we chose $l + 1 + \lambda < \log_2 N$, we have $z = x + r$. The distribution of \tilde{z} is statistically indistinguishable from the distribution of z (the distributions are distinguishable with an advantage of $2^{-\lambda}$ at most).

We also directly have that $(\text{SK}_P, \llbracket \tilde{z} \rrbracket) \equiv_s (\text{SK}_P, \llbracket z \rrbracket)$ and as a consequence, as the distribution of \tilde{z} and z is independent from \tilde{t}' and \tilde{r}_l ,

$$\begin{aligned} & (\text{PK}_{QR}, \text{SK}_P, \llbracket \tilde{z} \rrbracket; \text{coins}; [\tilde{t}'], [\tilde{r}_l]) \\ & \equiv_s (\text{PK}_{QR}, \text{SK}_P, \llbracket z \rrbracket; \text{coins}; [\tilde{t}'], [\tilde{r}_l]) \end{aligned}$$

By semantic security of QR,

$$\begin{aligned} & (\text{PK}_{QR}, \text{SK}_P, l, \llbracket z \rrbracket; \text{coins}; [\tilde{t}'], [\tilde{r}_l]) \\ & \equiv_c (\text{PK}_{QR}, \text{SK}_P, l, \llbracket z \rrbracket; \text{coins}; [t'], [r_l]) \end{aligned}$$

and

$$\begin{aligned} & S_B(\text{PK}_{QR}, \text{SK}_P, l) \\ & \equiv_c V_B(\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P) \end{aligned}$$

We conclude the proof of security using modular sequential composition. We replace the ideal calls for computing the encrypted bit $[t']$ by the provable secure DGK protocol and invoke Theorem B.2 to prove security in the semi-honest model. \square

Proof of Proposition A.2. The proof of security is similar to the one of Proposition A.1. Again we first suppose that $[t']$ is ideally computed (hybrid model).

A's view is $V_A = (\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{PK}_{QR}, \text{PK}_P; r, \text{coins}; [t'], [z_l])$ where PK_{QR} is the public key for the QR cryptosystem, PK_P is the public key for Paillier's cryptosystem and coins is the random tape used for the Paillier encryptions of r and 2^l , and the QR encryption of r_l .

Given $(\llbracket a \rrbracket, \llbracket b \rrbracket, \text{PK}_{QR}, \text{PK}_P)$, we build the simulator S_A :

1. Pick $\tilde{r} \leftarrow (0, 2^{\lambda+l}) \cap \mathbb{Z}$.
2. Generate $[\tilde{t}']$ and $[\tilde{z}_l]$, two encryptions of random bits under QR
3. Let $\widetilde{\text{coins}}$ be random coins for two Paillier encryptions and one QR encryption.
4. Output $(\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{PK}_{QR}, \text{PK}_P; \tilde{r}, \widetilde{\text{coins}}; [\tilde{z}_l])$

For both cases (A's view and the simulator S_A), r and \tilde{r} are taken from the same uniform distribution over $(0, 2^{\lambda+l}) \cap \mathbb{Z}$, and coins and $\widetilde{\text{coins}}$ are random tapes of the same length, so

$$\begin{aligned} & S_A(\llbracket a \rrbracket, \llbracket b \rrbracket, \text{PK}_{QR}, \text{PK}_P) \\ &= (\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{PK}_{QR}, \text{PK}_P; r, \text{coins}; [z_l]) \end{aligned}$$

By semantic security of the QR cryptosystem, we conclude with the computational indistinguishability of S_A and V_A distributions:

$$\begin{aligned} & S_A(\llbracket a \rrbracket, \llbracket b \rrbracket, \text{PK}_{QR}, \text{PK}_P) \\ &= (\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{PK}_{QR}, \text{PK}_P; r, \text{coins}; [z_l]) \\ & \equiv_c (\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{PK}_{QR}, \text{PK}_P; r, \text{coins}; [z_l]) \\ &= V_A(\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P) \end{aligned}$$

B's view is $V_B = (\text{SK}_{QR}, \text{SK}_P, \llbracket z \rrbracket, [t]; \text{coins})$ where SK_{QR} is the secret key for the QR cryptosystem, SK_P is the secret key for Paillier's cryptosystem, and coins are the random coins necessary for the QR encryption of z_l . The simulator $S_B(\text{SK}_{QR}, \text{SK}_P, a \leq b)$ runs as follows:

1. Compute $\widetilde{[t]}$ an encryption of the bit $(a \leq b)$ under QR.
2. Pick $\tilde{z} \leftarrow (0, 2^{\lambda+l}) \cap \mathbb{Z}$.
3. Encrypt \tilde{z} under Paillier: $\llbracket \tilde{z} \rrbracket$.
4. Let $\widetilde{\text{coins}}$ be random coins for one QR encryption.
5. Output $(\text{SK}_{QR}, \text{SK}_P, l, \llbracket \tilde{z} \rrbracket, [\tilde{t}]; \widetilde{\text{coins}})$

Once again, the distributions of coins and $\widetilde{\text{coins}}$ are identical:

$$\begin{aligned} & (\text{SK}_{QR}, \text{SK}_P, l, \llbracket \tilde{z} \rrbracket, [\tilde{t}]; \widetilde{\text{coins}}) \\ &= (\text{SK}_{QR}, \text{SK}_P, l, \llbracket \tilde{z} \rrbracket, [\tilde{t}]; \text{coins}) \end{aligned}$$

Recall that $z = x + r$ where x is an l bits integer and r is an $l + \lambda$ bits integer. The distribution of \tilde{z} is statistically indistinguishable from the distribution of z . We also directly have that $(\text{SK}_P, \llbracket \tilde{z} \rrbracket) \equiv_s (\text{SK}_P, \llbracket z \rrbracket)$ and as a consequence, as the distribution of \tilde{z} and z is independent from \tilde{t}' ,

$$\begin{aligned} & (\text{SK}_{QR}, \text{SK}_P, l, \llbracket \tilde{z} \rrbracket, [\tilde{t}]; \text{coins}) \\ & \equiv_s (\text{SK}_{QR}, \text{SK}_P, l, \llbracket z \rrbracket, [\tilde{t}]; \text{coins}) \end{aligned}$$

Moreover, by construction, $(\text{SK}_{QR}, [\tilde{t}]) = (\text{SK}_{QR}, [a < b])$ and

$$\begin{aligned} & (\text{SK}_{QR}, \text{SK}_P, l, \llbracket z \rrbracket, [\tilde{t}]; \text{coins}) \\ &= (\text{SK}_{QR}, \text{SK}_P, l, \llbracket z \rrbracket, [a < v]; \text{coins}). \end{aligned}$$

Finally, we have

$$\begin{aligned} & S_B(\text{SK}_{QR}, \text{SK}_P, a \leq b) \\ & \equiv_s V_B(\llbracket a \rrbracket, \llbracket b \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P). \end{aligned}$$

Again, we conclude the proof of security using modular sequential composition. We replace the ideal calls for computing the encrypted bit $[t']$ by the provable secure DGK protocol and invoke Theorem B.2 to prove security in the semi-honest model. \square

C.2 Argmax

Proof of Proposition 4.1. Correctness To prove correctness, we have to show that the following invariant holds: at the end of the loop for iteration i , m is the maximum of $\{a_{\pi(j)}\}_{1 \leq j \leq i}$ and $a_{\pi(i_0)} = m$.

If this holds, at the end of the loop iterations $a_{\pi(i_0)}$ is the maximum of $\{a_{\pi(j)}\}_{1 \leq j \leq k} = \{a_j\}_{1 \leq j \leq k}$, hence $i_0 = \text{argmax}_j a_{\pi(j)}$ and $\pi^{-1}(i_0) = \text{argmax}_j a_j$.

At initialization (line 4), the invariant trivially holds as the family $\{a_{\pi(j)}\}_{1 \leq j \leq i}$ contains only one element.

Suppose the property is true for iteration $i - 1$. Let us distinguish two cases:

- If b_i is true (i.e. $m \leq a_{\pi(i)}$), $\max\{a_{\pi(j)}\}_{1 \leq j \leq i-1} \leq a_{\pi(i)}$, as the invariant holds for the previous iteration, and then $\max\{a_{\pi(j)}\}_{1 \leq j \leq i} = a_{\pi(i)}$.

Then i_0 is set to i , $v_i = a'_i$ and $(x_i, y_i) = (0, 1)$. As a consequence, m is set by A to

$$v_i - x_i \cdot r_i - y_i \cdot s_i = a'_i - s_i = a_{\pi(i)}$$

We have clearly that $a_{\pi(i_0)} = a_{\pi(i)} = m$ and $m = \max\{a_{\pi(j)}\}_{1 \leq j \leq i}$, the invariant holds at the end of the i -th iteration in this case.

- If b_i is false ($m > a_{\pi(i)}$), $\max\{a_{\pi(j)}\}_{1 \leq j \leq i-1} > a_{\pi(i)}$ and $\max\{a_{\pi(j)}\}_{1 \leq j \leq i} = \max\{a_{\pi(j)}\}_{1 \leq j \leq i-1} = m$. Then i_0 is not changed, v_i is set to m'_i and $(x_i, y_i) = (1, 0)$. As a consequence,

$$v_i - x_i \cdot r_i - y_i \cdot s_i = m'_i - r_i = m$$

m is unchanged. As both m and i_0 stayed the same and $\max\{a_{\pi(j)}\}_{1 \leq j \leq i} = \max\{a_{\pi(j)}\}_{1 \leq j \leq i-1}$, the invariant holds at the end of the i -th iteration in this case.

Security We prove security in the hybrid model where line 5 of the protocol is ideally executed: we ask a trusted party T to compute the function $f(\llbracket x \rrbracket, \llbracket y \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P)$ in the f -hybrid model where

$$\begin{aligned} & f(\llbracket x \rrbracket, \llbracket y \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P) \\ &= (f_A(x, y, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P); \\ &= f_B(\llbracket x \rrbracket, \llbracket y \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P)) \end{aligned}$$

and f computes the function of Protocol 8, i.e. f_A returns nothing and f_B returns the bit $x \leq y$.

We will conclude using Theorem B.2.

A's view is

$$V_A = (\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P; \pi, \{r_i\}_{i=2}^k, \{s_i\}_{i=2}^k, \text{coins}; \{\llbracket v_i \rrbracket\}_{i=2}^k, \{(\llbracket x_i \rrbracket, \llbracket y_i \rrbracket)\}_{i=2}^k, \pi(\arg\max_i a_i))$$

where coins is the random tape for encryptions. To simulate A's real view, the simulator S_A does the following on input $(\llbracket a_1 \rrbracket, \dots, \llbracket a_k \rrbracket, l, \text{PK}_{QR}, \text{PK}_P, \arg\max_i a_i)$:

1. Picks a random permutation $\tilde{\pi}$ of $\{1, \dots, k\}$
2. Picks $k-1$ random integers $\tilde{r}_2, \dots, \tilde{r}_k$ in $(0, 2)^{l+\lambda} \cap \mathbb{Z}$
3. Picks $k-1$ random integers $\tilde{s}_2, \dots, \tilde{s}_k$ in $(0, 2)^{l+\lambda} \cap \mathbb{Z}$
4. Generates $k-1$ random Paillier encryptions $\llbracket \tilde{v}_2 \rrbracket, \dots, \llbracket \tilde{v}_k \rrbracket$.
5. Generates $k-1$ random bits \tilde{b}_i and encryptions $(\llbracket \tilde{x}_i \rrbracket, \llbracket \tilde{y}_i \rrbracket) = (\llbracket \tilde{b}_i \rrbracket, \llbracket \tilde{b}_i \rrbracket)$
6. Generate a random tape for $2(k-1)$ Paillier encryptions $\widetilde{\text{coins}}$
7. Outputs

$$(\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P; \tilde{\pi}, \{\tilde{r}_i\}_{i=2}^k, \{\tilde{s}_i\}_{i=2}^k, \widetilde{\text{coins}}; \{\llbracket \tilde{v}_i \rrbracket\}_{i=2}^k, \{(\llbracket \tilde{x}_i \rrbracket, \llbracket \tilde{y}_i \rrbracket)\}_{i=2}^k, \tilde{\pi}(\arg\max_i a_i))$$

We define the following hybrids:

- $H_0 = V_A(\llbracket a_1 \rrbracket, \dots, \llbracket a_k \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P)$
- $H_1 = (\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P; \pi, \{r_i\}_{i=2}^k, \{s_i\}_{i=2}^k, \text{coins}; \{\llbracket \tilde{v}_i \rrbracket\}_{i=2}^k, \{(\llbracket \tilde{x}_i \rrbracket, \llbracket \tilde{y}_i \rrbracket)\}_{i=2}^k, \pi(\arg\max_i a_i))$
- $H_2 = (\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P; \pi, \{\tilde{r}_i\}_{i=2}^k, \{\tilde{s}_i\}_{i=2}^k, \widetilde{\text{coins}}; \{\llbracket \tilde{v}_i \rrbracket\}_{i=2}^k, \{(\llbracket \tilde{x}_i \rrbracket, \llbracket \tilde{y}_i \rrbracket)\}_{i=2}^k, \pi(\arg\max_i a_i))$
- $H_3 = S_A(\llbracket a_1 \rrbracket, \dots, \llbracket a_k \rrbracket, l, \text{PK}_{QR}, \text{PK}_P, \arg\max_i a_i)$

By semantic security of Paillier's cryptosystem,

$$\begin{aligned} & (\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P; \pi, \{r_i\}_{i=2}^k, \{s_i\}_{i=2}^k; \\ & \quad \{\llbracket v_i \rrbracket\}_{i=2}^k, \{(\llbracket x_i \rrbracket, \llbracket y_i \rrbracket)\}_{i=2}^k, \pi(\arg\max_i a_i)) \\ & \equiv_c (\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P; \pi, \{\tilde{r}_i\}_{i=2}^k, \{\tilde{s}_i\}_{i=2}^k; \\ & \quad \{\llbracket \tilde{v}_i \rrbracket\}_{i=2}^k, \{(\llbracket \tilde{x}_i \rrbracket, \llbracket \tilde{y}_i \rrbracket)\}_{i=2}^k, \pi(\arg\max_i a_i)) \end{aligned}$$

and $H_0 \equiv_c H_1$ as $\pi(\arg\max_i a_i) = i_0$

Given that the \tilde{r}_i, \tilde{s}_i and $\widetilde{\text{coins}}$ are generated according to the same distribution as r_i, s_i (uniform over $(0, 2)^{l+\lambda} \cap \mathbb{Z}$) and coins (random tape for $2(k-1)$ Paillier encryptions), and that they are completely independent from the \tilde{v}_i or π , the hybrids H_1 and H_2 are equal.

Similarly, the distribution of $(\pi, \pi(\arg\max_i a_i))$ and $(\tilde{\pi}, \tilde{\pi}(\arg\max_i a_i))$ are exactly the same. As π and $\tilde{\pi}$ are independent from the other parameters, we also have $H_2 = H_3$. Hence, we showed that

$$\begin{aligned} & V_A(\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P) \\ & \equiv_c S_A(\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{PK}_{QR}, \text{PK}_P, \arg\max_i a_i). \end{aligned}$$

B's view is

$$V_B = (\text{SK}_P, \text{SK}_{QR}, l; \text{coins}; \{b_i\}_{i=2}^k, \{\llbracket m'_i \rrbracket\}_{i=2}^k, \{\llbracket a'_i \rrbracket\}_{i=2}^k)$$

where coins are the random coins for $k-1$ Paillier cyphertext refresh. The simulator $S_B(\text{SK}_P, \text{SK}_{QR}, l)$ runs as follows:

1. Generates a random permutation $\tilde{\pi}$ of $\{1, \dots, k\}$
2. Set $\llbracket \tilde{a}_i \rrbracket = \llbracket i \rrbracket$
3. Run the protocol with the $\llbracket \tilde{a}_i \rrbracket$ as input data, $\tilde{\pi}$ as the permutation, and same parameters otherwise. Let $(\text{SK}_P, \text{SK}_{QR}, l; \widetilde{\text{coins}}; \{b_i\}_{i=2}^k, \{\llbracket \tilde{m}'_i \rrbracket\}_{i=2}^k, \{\llbracket \tilde{a}'_i \rrbracket\}_{i=2}^k)$ be B's view of this run.
4. Outputs

$$(\text{SK}_P, \text{SK}_{QR}, l; \widetilde{\text{coins}}; \{b_i\}_{i=2}^k, \{\llbracket \tilde{m}'_i \rrbracket\}_{i=2}^k, \{\llbracket \tilde{a}'_i \rrbracket\}_{i=2}^k)$$

Let $p : \{a_i\}_{1 \leq i \leq k} \mapsto \{1, \dots, k\}$ be the function that associates a_i to its rank among the a_i (in ascendent order). Let us fix the permutation π for a while and define the following hybrids:

0. $H_0 = V_B(\{\llbracket a_i \rrbracket\}_{i=1}^k, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P)$
1. $H_1 = V_B(\{\llbracket p(a_1) \rrbracket\}_{i=1}^k, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P)$

We will show that these hybrids are statistically equal for every permutation π .

As $p(\cdot)$ is a map that does not change the order of the a_i , we have that for all i, j , $a_i \leq a_j \Leftrightarrow p(a_i) \leq p(a_j)$. As a consequence, for a given permutation π , the bits b_i do not change if we replace the a_i by $p(a_i)$. Similarly, the way the a'_i and m'_i are generated for H_0 and H_1 is the same: blinding by adding random noise from $(0, 2^{\lambda+l} \cap \mathbb{Z})$. Thus, $H_0 \equiv_s H_1$.

Now, we want to show that $H_1 \equiv_s S_B(\text{SK}_P, \text{SK}_{QR}, l)$ - we do not fix π anymore. Let π_0 be the permutation such that $p(a_i) = \pi_0(i)$. We can then rewrite H_1 as

$$H_1 = V_B(\llbracket \pi_0(1) \rrbracket, \dots, \llbracket \pi_0(k) \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P)$$

As $\tilde{\pi}$ and $\pi \circ \pi_0$ are statistically indistinguishable, we have $H_1 \equiv_s S_B(\text{SK}_P, \text{SK}_{QR}, l)$: recall that S_B 's output is the view of B when the protocol is run with the set $\{a_i = i\}$ as input set and $\tilde{\pi}$ as the permutation. Hence

$$\begin{aligned} & V_B(\llbracket a_1 \rrbracket, \dots, \llbracket a_k \rrbracket, l, \text{SK}_{QR}, \text{PK}_{QR}, \text{SK}_P, \text{PK}_P) \\ & \equiv_s S_B(\text{SK}_P, \text{SK}_{QR}, l) \end{aligned}$$

We conclude the proof of security using modular sequential composition. We replace the ideal calls for computing the encrypted bits b_i by the provable secure Protocol 8 and invoke Theorem B.2 to prove security in the semi-honest model. \square

C.3 Changing the encryption scheme

Proof of Proposition 4.2. In this protocol the computed function is probabilistic, and we have to show security according to the full definition (cf. section B.1). The function is f :

$$f(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2) = (\llbracket c \rrbracket_2, \emptyset)$$

For the sake of simplicity, we do not take into account the randomness used for the encryptions of r for A and c' for B. As before, the distribution of these coins for one party is completely independent of the other elements to be taken in account in the simulations, so we just do not mention them in security proof.

A's view is $V_A = (\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1; r; \llbracket c' \rrbracket_2)$. A's output is $\llbracket c \rrbracket_2$. The simulator $S_A(\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1)$ runs as follows:

1. Picks uniformly at random $\tilde{r} \leftarrow M$ and $\tilde{c}' \leftarrow M$.
2. Generates the encryption $\llbracket \tilde{c}' \rrbracket_2$ of \tilde{c}' under E_2 .
3. Outputs $(\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1; \tilde{r}; \llbracket \tilde{c}' \rrbracket_2)$.

r and \tilde{r} are taken from the same distribution, independently from any other parameter, so

$$\begin{aligned} & \{(\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1; \tilde{r}; \llbracket \tilde{c}' \rrbracket_2); f(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2)\} \\ & = \{(\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1; r; \llbracket \tilde{c}' \rrbracket_2); f(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2)\} \end{aligned}$$

(c' depends on r but does not appear in the previous distributions). By semantic security of scheme E_2 we have that

$$\begin{aligned} & \{(\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1; r; \llbracket \tilde{c}' \rrbracket_2); f(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2)\} \\ & \equiv_c \{(\text{PK}_1, \text{PK}_2, \llbracket c \rrbracket_1; r; \llbracket c' \rrbracket_2); \llbracket c \rrbracket_2\} \end{aligned}$$

and so

$$\begin{aligned} & \{S_A(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2), f(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2)\} \\ & \equiv_c \{V_A(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2), \text{Output}(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2)\} \end{aligned}$$

B's view is $V_B = (\text{SK}_1, \text{SK}_2; \llbracket c + r \rrbracket_1)$. We build a simulator $S_B(\text{SK}_1, \text{SK}_2)$:

1. Picks a random $\tilde{c} \leftarrow M$.
2. Encrypt \tilde{c} under E_1 .
3. Outputs $(\text{SK}_1, \text{SK}_2, \llbracket \tilde{c} \rrbracket_1)$.

Again, the distribution of \tilde{c} and $c + r$ are identical, so the real distribution $\{(\text{SK}_1, \text{SK}_2; \llbracket c + r \rrbracket_1); \llbracket c \rrbracket_2\}$ and the ideal distribution $\{(\text{SK}_1, \text{SK}_2; \llbracket \tilde{r} \rrbracket_1); f(\llbracket c \rrbracket_1, \text{PK}_1, \text{PK}_2, \text{SK}_1, \text{SK}_2)\}$ are statistically indistinguishable. \square

C.4 Computing dot products

Proof of Proposition 4.3. As B does not receive any message, its view only consists in its input and its random tape used for the encryptions. Hence the simulator S_B simply generate random coins and

$$S_B(y, \text{SK}_P) = (y, \text{SK}_P; \text{coins}) = V_B(x, y, \text{SK}_P, \text{PK}_P).$$

where rand are the random coins.

A's view is $V_A = (x, \text{PK}_P; r^A; \llbracket y_1 \rrbracket, \dots, \llbracket y_n \rrbracket)$. On input $(x, \text{PK}_P, \llbracket v \rrbracket)$, the simulator S_A does the following:

1. Generates n encryptions of 0 using Paillier: c_1, \dots, c_n .
2. Generates the random coins necessary for a Paillier re-randomization and put them in $\widetilde{\text{coins}}$.

3. Outputs $(x, \text{PK}_P; \widetilde{\text{coins}}; c_1, \dots, c_n)$.

coins and $\widetilde{\text{coins}}$ come from the same distribution, independently from other parameters. Thus,

$$\begin{aligned} & \{(x, \text{PK}_P; \widetilde{\text{coins}}; c_1, \dots, c_n); \llbracket \langle x, y \rangle \rrbracket\} \\ &= \{(x, \text{PK}_P; \text{coins}; c_1, \dots, c_n); \llbracket \langle x, y \rangle \rrbracket\} \end{aligned}$$

and by semantic security of Paillier,

$$\begin{aligned} & \{(x, \text{PK}_P; \text{coins}; c_1, \dots, c_n); \llbracket \langle x, y \rangle \rrbracket\} \\ & \equiv_c \{(x, \text{PK}_P; \text{coins}; \llbracket y_1 \rrbracket, \dots, \llbracket y_n \rrbracket\}; \llbracket v \rrbracket\} \end{aligned}$$

i.e., when f is $f(x, y, \text{SK}_P, \text{PK}_P) = (\llbracket \langle x, y \rangle \rrbracket, \emptyset)$

$$\begin{aligned} & \{S_A(x, \text{PK}_P, \llbracket v \rrbracket\}; f(x, y, \text{SK}_P, \text{PK}_P)\} \\ & \equiv_c \{V_A(x, y, \text{SK}_P, \text{PK}_P); \text{Output}(x, y, \text{SK}_P, \text{PK}_P)\} \end{aligned}$$

□

C.5 Classifiers

Hyperplane decision

Proof of Proposition 5.1. The client's view is

$$V_C = (\text{PK}_P, \text{PK}_{QR}, x; \{\llbracket v_i \rrbracket\}_{i=1}^k, i_0).$$

The simulator S_C , on input $(\text{PK}_P, \text{SK}_{QR}, x, k^*)$ where $k^* = \underset{i \in [k]}{\text{argmax}} \langle w_i, x \rangle$ does the following:

1. Generate k random Paillier encryptions $\llbracket \tilde{v}_i \rrbracket$
2. Output $(\text{PK}_P, \text{SK}_{QR}, x; \{\llbracket \tilde{v}_i \rrbracket\}_{i=1}^k, k^*)$

As the index i_0 that the client receives is its output, and as Paillier's cryptosystem is semantically secure, the distributions $S_C = (\text{PK}_P, \text{SK}_{QR}, x; \{\llbracket \tilde{v}_i \rrbracket\}_{i=1}^k, k^*)$ and $V_C = (\text{PK}_P, \text{SK}_{QR}, x; \{\llbracket v_i \rrbracket\}_{i=1}^k, i_0)$ are computationally indistinguishable.

As the server views nothing but its inputs (the server does not receive any message in the hybrid model), we use for the trivial simulator that just outputs its inputs for the proof of security.

As Protocols 1 and 3 are secure in the honest-but-curious model, we obtain the security of the hyperplane decision protocol using modular sequential composition (Theorem B.2). □

Bayes classifier

Proof of Proposition 6.1. The client's view is

$$V_C = (\text{PK}_P, \text{SK}_{QR}, x; \llbracket P \rrbracket, \{\llbracket T_{i,j} \rrbracket\}, i_0).$$

The simulator S_C , on input $(\text{PK}_P, \text{SK}_{QR}, x, i_{max})$ where $i_{max} = \underset{j}{\text{argmax}} \mathbb{P}(C = c_j | X = x)$,

- generates tables of random Paillier encryptions $\llbracket \tilde{P} \rrbracket$ and $\{\llbracket T_{i,j} \rrbracket\}$;
- outputs $(\text{PK}_P, \text{SK}_{QR}, x; \llbracket \tilde{P} \rrbracket, \{\llbracket \tilde{T}_{i,j} \rrbracket\}, i_{max})$.

As the integer i_0 that the client receives is its output, and as Paillier's cryptosystem is semantically secure, the distributions $S_C = (\text{PK}_P, \text{SK}_{QR}, x; \llbracket \tilde{P} \rrbracket, \{\llbracket \tilde{T}_{i,j} \rrbracket\}, i_{max})$ and $V_C = (\text{PK}_P, \text{SK}_{QR}, x; \llbracket P \rrbracket, \{\llbracket T_{i,j} \rrbracket\}, i_0)$ are computationally indistinguishable.

Again, as the server views nothing but its inputs (the server does not receive any message in the hybrid model), we use the trivial simulator that outputs its inputs and the random coins for the encryption for the proof of security.

As Protocol 1 is secure in the honest-but-curious model, we obtain the security of the hyperplane decision protocol using modular sequential composition (Theorem B.2). \square

Decision tree

Proof of Proposition 7.1. The proof of security for the server is very easily obtained using modular sequential composition of the comparison protocol and Protocol 2: in the hybrid model, the client receives nothing but the encrypted result.

For the client also the proof is trivial, using modular sequential composition and the semantical security of QR and of the FHE scheme: the encryptions of bits b_i are computational indistinguishable from random bits whether they are encrypted under QR or the FHE scheme. \square