PINFER: Privacy-Preserving Inference

Logistic Regression, Support Vector Machines, and More, Over Encrypted Data

Marc Joye and Fabien Petitcolas

OneSpan, Brussels, Belgium
{marc.joye,fabien.petitcolas}@onespan.com

Abstract. The foreseen growing role of outsourced machine learning services is raising concerns about the privacy of user data. This paper proposes a variety of protocols for privacy-preserving regression and classification that (i) only require additively homomorphic encryption algorithms, (ii) limit interactions to a mere request and response, and (iii) that can be used directly for important machine-learning algorithms such as logistic regression and SVM classification. The basic protocols are then extended and applied to simple feed-forward neural networks.

Keywords: Machine learning as a service · Linear regression · Logistic regression · Support vector machines · Feed-forward neural networks · Data privacy · Additively homomorphic encryption

1 Introduction

The popularity and hype around machine learning, combined with the explosive growth of user-generated data, is pushing the development of *machine learning as a service* (MLaaS). A typical application scenario of MLaaS is shown in Fig. 1. It involves a client sending data to a service provider (server) owning and running a trained machine learning model for a given task (e.g., medical diagnosis). Both the input data and the model should be kept private: for obvious privacy reasons on the client's side and to protect intellectual property on the server's side.

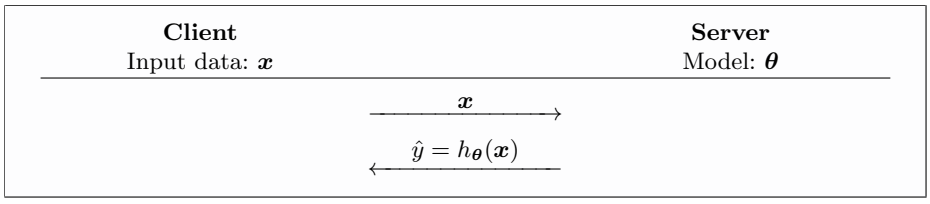

Fig. 1: A server offering MLaaS owns a model defined by its parameters $\boldsymbol{\theta}$. A client needs the prediction $h_{\boldsymbol{\theta}}(\boldsymbol{x})$ of this model for a new input data $\boldsymbol{x}$. This prediction is a function of the model and of the data.

In this paper we look at various protocols allowing the realisation of such scenario in a minimum number of message exchanges between both parties. Our assumption is that both the client and the server are honest but curious, that is, they both follow the protocol but may record information all along with the aim, respectively, to learn the model and to breach the client's privacy. Our design is guided by the following *ideal* requirements, in decreasing importance:

1. **Input confidentiality**—The server does not learn anything about the input data $\boldsymbol{x}$ provided by the client;
2. **Output confidentiality**—The server does not learn the outcome $\hat{y}$ of the calculation;
3. **Minimal model leakage**—The client does not learn any other information about the model beyond what is revealed by the successive outputs.

With respect to the issue of model leakage, it is noted that the client gets access to the outcome, i.e., the value of $h_{\boldsymbol{\theta}}(\boldsymbol{x})$, which may leak information about $\boldsymbol{\theta}$, violating Requirement 3. This is unavoidable and not considered as an attack within our framework. Possible countermeasures to limit the leakage on the model include rounding the output or adding some noise to it [19].

Related Work Earliest works for private machine learning evaluation [2,17] were concerned with training models in a privacy-preserving manner. More recent implementations for linear regression, logistic regression, as well as neural networks are offered by SecureML [18]. The case of support vector machines (SVM) is, for example, covered in [21]. On the contrary, this paper deals with the problem of privately *evaluating* a linear machine-learning model, including linear/logistic regression and SVM classification. In [4], Bos et al. suggest to evaluate a logistic regression model by replacing the sigmoid function with its Taylor series expansion. They then apply fully homomorphic encryption so as to get the output result through a series of multiplications and additions over encrypted data. They observe that using terms up to degree 7 the Taylor expansion gives roughly two digits of accuracy to the right decimal. Kim et al. [16] argue that such an expansion does not provide enough accuracy on real-world data sets and propose another polynomial approximation. For SVM classification, Zhang et al. [21, Protocol 2] propose to return an encryption of the raw output. The client decrypts it and applies the discriminating function to obtain the corresponding class. Unfortunately, this leaks more information than necessary on the model. A similar path is taken by Barni et al. in [3] for feed-forward neural networks. Extracting the model (even partially) is nevertheless more difficult in their case because of the inherent complexity of the model. Moreover, to further obfuscate it (and thereby limit the potential leakage), the authors suggest to randomly permute computing units (neurons) sharing the same activation function or to add dummy ones. For classification, the approach put forward by Bost et al. [5] is closest to ours. They construct three classification protocols fulfilling our design criteria (Requirements 1–3): hyperplane decision, naïve Bayes, and decision trees. An approach orthogonal to ours that introduces

privacy in regression or classification is differential privacy [8]. Crucially, it can be combined with secure computation, in our case by incorporating noise in the input vectors or in the model parameters. Differential privacy can thus be used to enhance the privacy properties of our protocols.

Our Contributions Our paper follows the line of work by Bost et al., making use only of *additively* homomorphic encryption (i.e., homomorphic encryption supporting additions). We devise new privacy-preserving protocols for a variety of important prediction tasks. The protocols we propose either improve on [5] or address machine-learning models not covered in [5]. In particular, we aim at minimising the number of message exchanges to a mere request and response. This is important when latency is critical. An application of [5, Protocol 4] to binary SVM classification adds a round-trip to the comparison protocol whereas our implementation optimally only needs a *single* round-trip, all included. Likewise, a single round-trip is needed in our private logistic regression protocol, as in [4,21]. But contrary to [4,21], the resulting prediction is *exact* in our case (i.e., there is no loss of accuracy) and does not require the power of fully homomorphic encryption. With respect to neural networks, we adapt our protocols to binarised networks and to networks relying of the popular ReLU activation; see Section 4.2. As far as we know, this results in the first privacy-preserving implementation of the non-linear ReLU function from additively homomorphic encryption.

2 Preliminaries

2.1 Linear Models and Beyond

Owing to their simplicity, linear models (see, e.g., [1, Chapter 3] or [13, Chapters 3 and 4]) should not be overlooked: They are powerful tools for a variety of machine learning tasks and find numerous applications.

Problem Setup In a nutshell, machine learning works as follows. Each particular problem instance is characterised by a set of d features which may be viewed as a vector $(x_1, \ldots, x_d)^\intercal$ of $\mathbb{R}^d$. For practical reasons, a fixed coordinate $x_0 = 1$ is added. We let $\mathcal{X} \subseteq \{1\} \times \mathbb{R}^d$ denote the input space and $\mathcal{Y}$ the output space.

There are two phases:

- The *learning phase* consists in approximating a target function $f\colon \mathcal{X} \to \mathcal{Y}$ from $\mathcal{D} = \big\{(\boldsymbol{x_i}, y_i) \in \mathcal{X} \times \mathcal{Y} \mid y_i = f(\boldsymbol{x_i})\big\}_{1 \leqslant i \leqslant n}$, a training set of n pairs of elements. Note that the target function can be noisy. The output of the learning phase is a function $h_{\boldsymbol{\theta}}\colon \mathcal{X} \to \mathcal{Y}$ drawn from some hypothesis set of functions.
- In the *testing phase*, when a new data point $\boldsymbol{x} \in \mathcal{X}$ comes in, it is evaluated on $h_{\boldsymbol{\theta}}$ as $\hat{y} = h_{\boldsymbol{\theta}}(\boldsymbol{x})$. The hat on y indicates that it is a predicted value.

Since $h_{\boldsymbol{\theta}}$ was chosen in a way to "best match" f on the training set $\mathcal{D}$, it is expected that it will provide a good approximation on a new data point. Namely, we have $h_{\boldsymbol{\theta}}(\boldsymbol{x_i}) \approx y_i$ for all $(\boldsymbol{x_i}, y_i) \in \mathcal{D}$ and we should have $h_{\boldsymbol{\theta}}(\boldsymbol{x}) \approx f(\boldsymbol{x})$ for $(\boldsymbol{x}, \cdot) \notin \mathcal{D}$. Of course, this highly depends on the problem under consideration, the data points, and the hypothesis set of functions.

In particular, linear models for machine learning use a hypothesis set of functions of the form $h_{\boldsymbol{\theta}}(\boldsymbol{x}) = g(\boldsymbol{\theta}^\intercal \boldsymbol{x})$ where $\boldsymbol{\theta} = (\theta_0, \theta_1, \ldots, \theta_d)^\intercal \in \mathbb{R}^{d+1}$ are the model parameters and $g\colon \mathbb{R} \to \mathcal{Y}$ is a function mapping the linear calculation to the output space. Specific choices for g are discussed hereafter.

Linear Regression A *linear regression* model assumes that the real-valued target function f is linear—or more generally affine—in the input variables. In other words, it is based on the premise that f is well approximated by an affine map; i.e., g is the identity map: $f(\boldsymbol{x_i}) \approx g(\boldsymbol{\theta}^\intercal \boldsymbol{x_i}) = \boldsymbol{\theta}^\intercal \boldsymbol{x_i}$, $1 \leqslant i \leqslant n$, for training data $\boldsymbol{x_i} \in \mathcal{X}$ and weight vector $\boldsymbol{\theta} \in \mathbb{R}^{d+1}$.

The linear regression algorithm relies on the least squares method to find the coefficients of $\boldsymbol{\theta}$. Once $\boldsymbol{\theta}$ has been computed, it can be used to produce estimates on new data points $\boldsymbol{x} \in \mathcal{X}$ as $\hat{y} = \boldsymbol{\theta}^\intercal \boldsymbol{x}$.

Support Vector Machines Another important problem is to classify data into different classes. This corresponds to a target function f whose range $\mathcal{Y}$ is discrete. Of particular interest is the case of two classes, say $+1$ and -1, in which case $\mathcal{Y} = \{+1, -1\}$.

In dimension d, an hyperplane $\mathit{\Pi}$ is given by an equation of the form $\theta_0 + \theta_1 X_1 + \theta_2 X_2 + \cdots + \theta_d X_d = 0$ where $\boldsymbol{\theta'} = (\theta_1, \ldots, \theta_d)^\intercal$ is the normal vector to $\mathit{\Pi}$ and $\theta_0 / \|\boldsymbol{\theta'}\|$ indicates the offset from the origin. When the training data are *linearly separable*, there is some hyperplane $\mathit{\Pi}$ such that for each $(\boldsymbol{x_i}, y_i) \in \mathcal{D}$, one has $y_i\, \boldsymbol{\theta}^\intercal \boldsymbol{x_i} \geqslant 1$, $(1 \leqslant i \leqslant n)$. The training data points $\boldsymbol{x_i}$ satisfying $y_i\, \boldsymbol{\theta}^\intercal \boldsymbol{x_i} = 1$ are called *support vectors*. When the training data are not linearly separable, it is not possible to satisfy the previous hard constraint $y_i\, \boldsymbol{\theta}^\intercal \boldsymbol{x_i} \geqslant 1$, $(1 \leqslant i \leqslant n)$. So-called "slack variables" $\xi_i = \max(0, 1 - y_i\, \boldsymbol{\theta}^\intercal \boldsymbol{x_i})$ are generally introduced in the optimisation problem. They tell how large a violation of the hard constraint there is on each training point.

There are many possible choices for $\boldsymbol{\theta}$. For better classification, the separating hyperplane $\mathit{\Pi}$ is chosen so as to maximise the *margin*; namely, the minimal distance between any training data point and $\mathit{\Pi}$. Now, from the resulting model $\boldsymbol{\theta}$, when a new data point $\boldsymbol{x}$ comes in, its class is estimated as the sign of the discriminating function $\boldsymbol{\theta}^\intercal \boldsymbol{x}$; i.e., $\hat{y} = \operatorname{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x})$.

When there are more than two classes, the optimisation problem returns several vectors $\boldsymbol{\theta_k}$, each defining a boundary between a particular class and all the others. The classification problem becomes an iteration to find out which $\boldsymbol{\theta_k}$ maximises $\boldsymbol{\theta_k}^\intercal \boldsymbol{x}$ for a given test point $\boldsymbol{x}$.

Logistic Regression *Logistic regression* is widely used in predictive analysis to output a probability of occurrence. The logistic function is defined by the

sigmoid function $\sigma\colon \mathbb{R} \to [0,1],\ t \mapsto \sigma(t) = \frac{1}{1+e^{-t}}$. The logistic regression model returns $h_{\boldsymbol{\theta}}(\boldsymbol{x}) = \sigma(\boldsymbol{\theta}^\intercal \boldsymbol{x}) \in [0,1]$, which can be interpreted as the probability that $\boldsymbol{x}$ belongs to the class $y = +1$. The SVM classifier thresholds the value of $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ around 0, assigning to $\boldsymbol{x}$ the class $y = +1$ if $\boldsymbol{\theta}^\intercal \boldsymbol{x} > 0$ and the class $y = -1$ if $\boldsymbol{\theta}^\intercal \boldsymbol{x} < 0$. In this respect, the logistic function is seen as a soft threshold as opposed to the hard threshold, $+1$ or -1, offered by SVM. Other threshold functions are possible. Another popular soft threshold relies on tanh, the hyperbolic tangent function, whose output range is $[-1,1]$.

Remark 1. Because the logistic regression algorithm predicts probabilities rather than just classes, we fit it through likelihood optimisation. Specifically, given the training set $\mathcal{D}$, we learn the model by maximising $\prod_{y_i=+1} p_i \cdot \prod_{y_i=-1}(1-p_i)$ where $p_i = \sigma(\boldsymbol{\theta}^\intercal \boldsymbol{x_i})$. This deviates from the general description of our problem setup, where the learning is directly done on the pairs $(\boldsymbol{x_i}, y_i)$. However, the testing phase is unchanged: the outcome is expressed as $h_{\boldsymbol{\theta}}(\boldsymbol{x}) = \sigma(\boldsymbol{\theta}^\intercal \boldsymbol{x})$. It therefore fits our framework for private inference, that is, the private evaluation of $h_{\boldsymbol{\theta}}(\boldsymbol{x}) = g(\boldsymbol{\theta}^\intercal \boldsymbol{x})$ for a certain function g; the sigmoid function σ in this case.

2.2 Cryptographic Tools

Representing Real Numbers So far, we have discussed a number of machine learning models using real numbers, but the cryptographic tools we intend to use require working on integers. In order to operate over encrypted data, we need to accurately represent real numbers as elements of the finite message set $\mathcal{M} \subset \mathbb{Z}$.

To do that, since all input variables of machine learning models are typically rescaled in the range $[-1,1]$, one could use a fixed point representation. A real number x with a fractional part of at most P bits uniquely corresponds to signed integer $z = x \cdot 2^P$. Hence, with a fixed-point representation, a real number x is represented by $z = \lfloor x \cdot 2^P \rfloor$, where integer P is called the bit-precision. The sum of $x_1, x_2 \in \mathbb{R}$ is performed as $z_1 + z_2$ and their multiplication as $\lfloor (z_1 \cdot z_2)/2^P \rfloor$.

Additively Homomorphic Encryption It is useful to introduce some notation. We let $[\![\cdot]\!]$ and $]\!]\cdot[\!]$ denote the encryption and decryption algorithms, respectively. The message space is an additive group $\mathcal{M} \cong \mathbb{Z}/M\mathbb{Z}$. It consists of integers modulo M and we view it as $\mathcal{M} = \{-\lfloor M/2 \rfloor, \ldots, \lceil M/2 \rceil - 1\}$ in order to keep track of the sign. Ciphertexts are noted with Gothic letters. If $\boldsymbol{m} = (m_1, \ldots, m_d) \in \mathcal{M}^d$ is a vector, we write $\mathfrak{m} = [\![\boldsymbol{m}]\!]$ as a shorthand for $(\mathfrak{m}_1, \ldots, \mathfrak{m}_d) = ([\![m_1]\!], \ldots, [\![m_d]\!])$. The minimal security notion that we require is semantic security [12]; in particular, encryption is probabilistic.

Algorithm $[\![\cdot]\!]$ being *additively* homomorphic (over $\mathcal{M}$) means that given any two plaintext messages m_1 and m_2 and their corresponding ciphertexts $\mathfrak{m}_1 = [\![m_1]\!]$ and $\mathfrak{m}_2 = [\![m_2]\!]$, we have $\mathfrak{m}_1 \boxplus \mathfrak{m}_2 = [\![m_1 + m_2]\!]$ and $\mathfrak{m}_1 \boxminus \mathfrak{m}_2 = [\![m_1 - m_2]\!]$ for some publicly known operations $\boxplus$ and $\boxminus$ on ciphertexts. By induction, for a given integer scalar $r \in \mathbb{Z}$, we also have $[\![r \cdot m_1]\!] \coloneqq r \odot \mathfrak{m}_1$.

It is worth noting here that the decryption of $(\mathfrak{m}_1 \boxplus \mathfrak{m}_2)$ gives $(m_1 + m_2)$ as an element of $\mathcal{M}$; that is, $\llbracket \mathfrak{m}_1 \boxplus \mathfrak{m}_2 \rrbracket \equiv m_1 + m_2 \pmod{M}$. Similarly, we also have $\llbracket \mathfrak{m}_1 \boxminus \mathfrak{m}_2 \rrbracket \equiv m_1 - m_2 \pmod{M}$ and $\llbracket r \odot \mathfrak{m}_1 \rrbracket \equiv r \cdot m_1 \pmod{M}$.

Private Comparison Protocol In [6,7], Damgård et al. present a protocol for comparing private values. It was later extended and improved in [9] and [20,15]. The protocol makes use of an additively homomorphic encryption scheme. It compares two non-negative ℓ-bit integers. The message space is $\mathcal{M} \cong \mathbb{Z}/M\mathbb{Z}$ with $M \geqslant 2^\ell$ and is supposed to behave like an integral domain.

DGK+ protocol A client possesses a private ℓ-bit value $\mu = \sum_{i=0}^{\ell-1} \mu_i \, 2^i$ and a server possesses a private ℓ-bit value $\eta = \sum_{i=0}^{\ell-1} \eta_i \, 2^i$. They seek to respectively obtain bits δ_C and δ_S such that $\delta_C \oplus \delta_S = [\mu \leqslant \eta]$ (where $\oplus$ represents the exclusive OR operator, and $[\mathsf{Pred}] = 1$ if predicate Pred is true, and 0 otherwise). Following [15, Fig. 1], the DGK+ protocol proceeds in four steps:

1. The client encrypts each bit μ_i of μ under its public key and sends $\llbracket \mu_i \rrbracket$, $0 \leqslant i \leqslant \ell - 1$, to the server.
2. The server chooses uniformly at random a bit δ_S and defines $s = 1 - 2\delta_S$. It also selects $\ell + 1$ random non-zero scalars $r_i \in \mathcal{M}$, $-1 \leqslant i \leqslant \ell - 1$.
3. Next, the server computes[1]

$$\begin{cases} \llbracket h_i^* \rrbracket = r_i \odot \left(\llbracket 1 \rrbracket \boxplus \llbracket s \cdot \mu_i \rrbracket \boxminus \llbracket s \cdot \eta_i \rrbracket \boxplus \left(\boxplus_{j=i+1}^{\ell-1} \llbracket \mu_j \oplus \eta_j \rrbracket \right) \right) \\ \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \text{for } \ell - 1 \geqslant i \geqslant 0\,, \\ \llbracket h_{-1}^* \rrbracket = r_{-1} \odot \left(\llbracket \delta_S \rrbracket \boxplus \boxplus_{j=0}^{\ell-1} \llbracket \mu_j \oplus \eta_j \rrbracket \right) \end{cases} \quad (1)$$

   and sends the $\ell + 1$ ciphertexts $\llbracket h_i^* \rrbracket$ in a *random order* to the client.
4. Using its private key, the client decrypts the received $\llbracket h_i^* \rrbracket$'s. If one is decrypted to zero, the client sets $\delta_C = 1$. Otherwise, it sets $\delta_C = 0$.

Remark 2. At this point, neither the client, nor the server, knows whether $\mu \leqslant \eta$ holds. One of them (or both) needs to reveal its share of δ $(= \delta_C \oplus \delta_S)$ so that the other can find out. Following the original DGK protocol [6], this modified comparison protocol is secure in the semi-honest model (i.e., against honest but curious adversaries).

3 Basic Protocols of Privacy-Preserving Inference

In this section, we present three families of protocols for private inference. They aim to satisfy the ideal requirements given in the introduction while keeping the number of exchanges to a bare minimum. Interestingly, they only make use of additively homomorphic encryption.

We keep the general model presented in the introduction, but now work with integers only. The client holds $\boldsymbol{x} = (1, x_1, \ldots, x_d)^\intercal \in \mathcal{M}^{d+1}$, a private feature

[1] Given $\llbracket \mu_i \rrbracket$, the server obtains $\llbracket \eta_i \oplus \mu_i \rrbracket$ as $\llbracket \mu_i \rrbracket$ if $\eta_i = 0$, and as $\llbracket 1 \rrbracket \boxminus \llbracket \mu_i \rrbracket$ if $\eta_i = 1$.

vector, and the server possesses a trained machine-learning model given by its parameter vector $\boldsymbol{\theta} = (\theta_0, \ldots, \theta_d)^\intercal \in \mathcal{M}^{d+1}$ or, in the case of feed-forward neural networks a set of matrices made of such vectors. At the end of protocol, the client obtains the value of $g(\boldsymbol{\theta}^\intercal \boldsymbol{x})$ for some function g and learns nothing else; the server learns nothing. To make the protocols easier to read, for a real-valued function g, we abuse notation and write $g(t)$ for an integer t assuming g also includes the conversion to real values; see Section 2.2. We also make the distinction between the encryption algorithm $[\![\cdot]\!]$ using the client's public key and the encryption algorithm $\{\!\{\cdot\}\!\}_s$ using the server's public key and stress that, not only keys are different, but the algorithm could also be different. We use $]\!]\cdot[\![$ and $\}\!\}\cdot\{\!\{_s$ for the respective corresponding decryption algorithms.

3.1 Private Linear/Logistic Regression

Evaluating a linear regression model over encrypted data is straightforward and requires just one round of communication; see Appendix A.1. Things get more complicated for logistic regression. At first sight, it seems counter-intuitive that additively homomorphic encryption could suffice to evaluate a logistic regression model over encrypted data. After all, the sigmoid function, $\sigma(t)$, is non-linear. The key observation is that the sigmoid function is *injective*. So the client does not learn more about the model $\boldsymbol{\theta}$ from $t \coloneqq \boldsymbol{\theta}^\intercal \boldsymbol{x}$ than it can learn from $\hat{y} \coloneqq \sigma(t)$ since the value of t can be recovered from $\hat{y}$ using $t = \sigma^{-1}(\hat{y}) = \ln(\hat{y}/(1-\hat{y}))$.

Our Core Protocol The protocol we propose for privacy-preserving linear or logistic regression is detailed in Fig. 2. Let (pk_C, sk_C) denote the client's matching pair of public encryption key/private decryption key for an additively homomorphic encryption scheme $[\![\cdot]\!]$. We use the notation of Section 2.2. If B is an upper bound on the inner product (in absolute value), the message space $\mathcal{M} = \{-\lfloor M/2 \rfloor, \ldots, \lceil M/2 \rceil - 1\}$ should be such that $M \geqslant 2B + 1$.

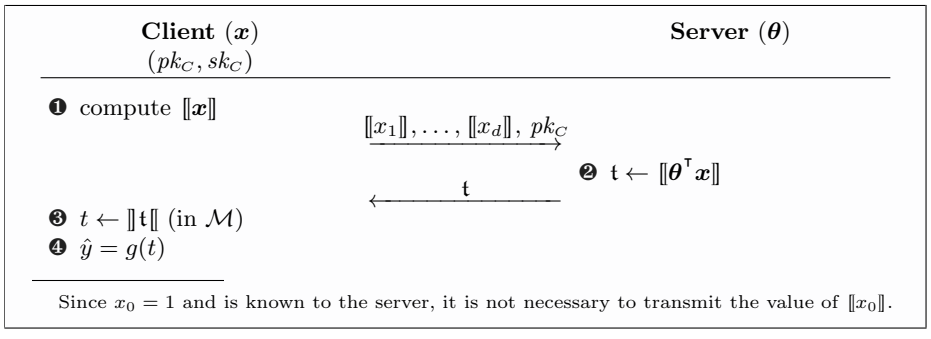

Fig. 2: Privacy-preserving regression. Encryption is done using the client's public key and noted $[\![\cdot]\!]$. The server learns nothing. Function g is the identity map for linear regression and the sigmoid function for logistic regression.

1. In a first step, the client encrypts its feature vector $\boldsymbol{x} \in \mathcal{M}^{d+1}$ under its public key pk_C and gets $[\![\boldsymbol{x}]\!] = ([\![x_0]\!], [\![x_1]\!], \ldots, [\![x_d]\!])$. The ciphertext $[\![\boldsymbol{x}]\!]$ along with the client's public key are sent to the server.
2. In a second step, from its model $\boldsymbol{\theta}$, the server computes an encryption of the inner product over encrypted data as: $\mathfrak{t} = [\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!] = [\![\theta_0]\!] \boxplus \boxplus_{j=1}^{d} \theta_j \odot [\![x_j]\!]$. The server returns $\mathfrak{t}$ to the client.
3. In a third step, the client uses its private decryption key sk_C to decrypt $\mathfrak{t}$, and gets the inner product $t = \boldsymbol{\theta}^\intercal \boldsymbol{x}$ as a signed integer of $\mathcal{M}$.
4. In a final step, the client applies the g function to obtain the prediction $\hat{y}$ corresponding to input vector $\boldsymbol{x}$.

Dual Approach The previous protocol encrypts with the client's public key pk_C. In the dual approach, the server's public key is used for encryption. Let (pk_S, sk_S) denote the public/private key pair of the server for some additively homomorphic encryption scheme $\left(\{\!|\cdot|\!\}_s, |\!\{\cdot\}\!|_s\right)$. The message space $\mathcal{M}$ is unchanged.

In this case, the server needs to publish an encrypted version $\{\!|\boldsymbol{\theta}|\!\}_s$ of its model. The client must therefore get a copy of $\{\!|\boldsymbol{\theta}|\!\}_s$ once, but can then engage in the protocol as many times as it wishes. One could also suppose that each client receives a different encryption of $\boldsymbol{\theta}$ using a server's encryption key specific to the client, or that a key rotation is performed on a regular basis. This protocol uses a mask μ which is chosen uniformly at random in $\mathcal{M}$. Consequently, it is important to see that t^* ($\equiv \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu \pmod{M}$) is also uniformly distributed over $\mathcal{M}$. Thus, the server gains no bit of information from t^*. The different steps are summarised in Fig. 3.

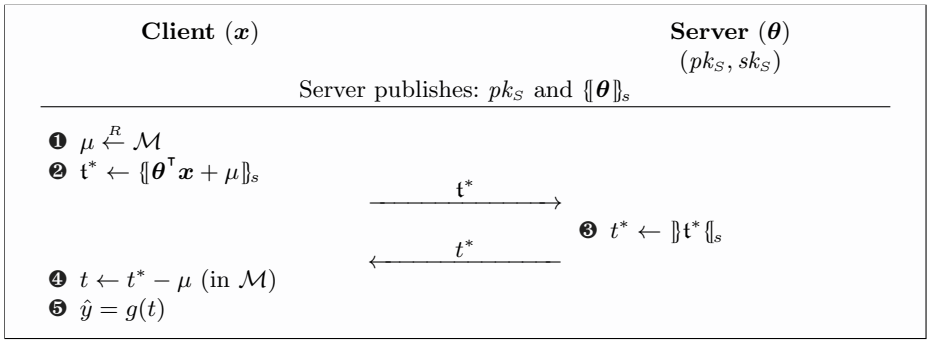

Fig. 3: Dual approach for privacy-preserving regression. Here, encryption is done using the server's public key pk_S and noted $\{\!|\cdot|\!\}_s$. Function g is the identity map for linear regression and the sigmoid function for logistic regression.

Extensions The proposed methods are not limited to the identity map or the sigmoid function but generalise to any injective function g. This includes

the hyperbolic tangent function, the arctangent function, the softsign function, the softplus function, the leaky ReLU function, and more. For any injective function g, there is no more information leakage in returning $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ than $g(\boldsymbol{\theta}^\intercal \boldsymbol{x})$.

3.2 Private SVM Classification

As discussed in Section 2.1, SVM inference can be abridged to the evaluation of the sign of an inner product. However, the sign function is clearly not injective. Our idea is to make use of a privacy-preserving comparison protocol. For concreteness, we consider the DGK+ protocol; but any privacy-preserving comparison protocol could be adapted.

Our Core Protocol As explained in Appendix A.2, a naïve implementation for private SVM has important issues. To address them we select the message space much larger than the upper bound B on the inner product.

Specifically, if $\boldsymbol{\theta}^\intercal \boldsymbol{x} \in [-B, B]$ then, letting ℓ be the bit-length of B, the message space $\mathcal{M} = \{-\lfloor M/2 \rfloor, \ldots, \lceil M/2 \rceil - 1\}$ should be dimensioned such that $M \geqslant 2^\ell(2^\kappa + 1) - 1$ for some security parameter κ. Let μ be an $(\ell + \kappa)$-bit integer that is chosen such that $\mu \geqslant B$. By construction we will then have $0 \leqslant \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu < M$ so that the decrypted value modulo M corresponds to the actual integer value. As will become apparent, this presents the further advantage of optimising the bandwidth requirements: the number of exchanged ciphertexts depends on the length of B and not on the length of M (notice that $M = \#\mathcal{M}$).

Our resulting core protocol for private SVM classification of a feature vector $\boldsymbol{x}$ is illustrated in Fig. 4 and includes the following steps:

0. The server publishes pk_S and $\{\!|\boldsymbol{\theta}|\!\}_s$.
1. Let κ be a security parameter. The client starts by picking uniformly at random in $[2^\ell - 1, 2^{\ell+\kappa})$ an integer $\mu = \sum_{i=0}^{\ell+\kappa-1} \mu_i \, 2^i$.
2. In a second step, the client computes, over encrypted data, the inner product $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ and masks the result with μ to get $\mathfrak{t}^* = \{\!|t^*|\!\}_s$ with $t^* = \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$.
3. Next, the client individually encrypts the first ℓ bits of μ with its own encryption key to get $[\![\mu_i]\!]$, for $0 \leqslant i \leqslant \ell - 1$, and sends $\mathfrak{t}^*$ and the $[\![\mu_i]\!]$'s to the server.
4. Upon reception, the server decrypts $\mathfrak{t}^*$ to get $t^* := |\!|\mathfrak{t}^*|\!|_s \bmod M = \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$ and defines the ℓ-bit integer $\eta := t^* \bmod 2^\ell$.
5. The DGK+ protocol is now applied to two ℓ-bit values $\underline{\mu} := \mu \bmod 2^\ell = \sum_{i=0}^{\ell-1} \mu_i \, 2^i$ and $\eta = \sum_{i=0}^{\ell-1} \eta_i \, 2^i$. The server selects the $(\ell+1)$-th bit of t^* for δ_S (i.e., $\delta_S = \lfloor t^*/2^\ell \rfloor \bmod 2$), defines $s = 1 - 2\delta_S$, and forms the $[\![h_i^*]\!]$'s (with $-1 \leqslant i \leqslant \ell - 1$) as defined by Eq. (1). The server permutes randomly the $[\![h_i^*]\!]$'s and sends them to the client.
6. The client decrypts the $[\![h_i^*]\!]$'s and gets the h_i^*'s. If one of them is zero, it sets $\delta_C = 1$; otherwise it sets $\delta_C = 0$.
7. As a final step, the client obtains the predicted class as $\hat{y} = (-1)^{\neg(\delta_C \oplus \mu_\ell)}$, where μ_ℓ denotes bit number ℓ of μ.

Again, the proposed protocol keeps the number of interactions between the client and the server to a minimum: a request and a response.

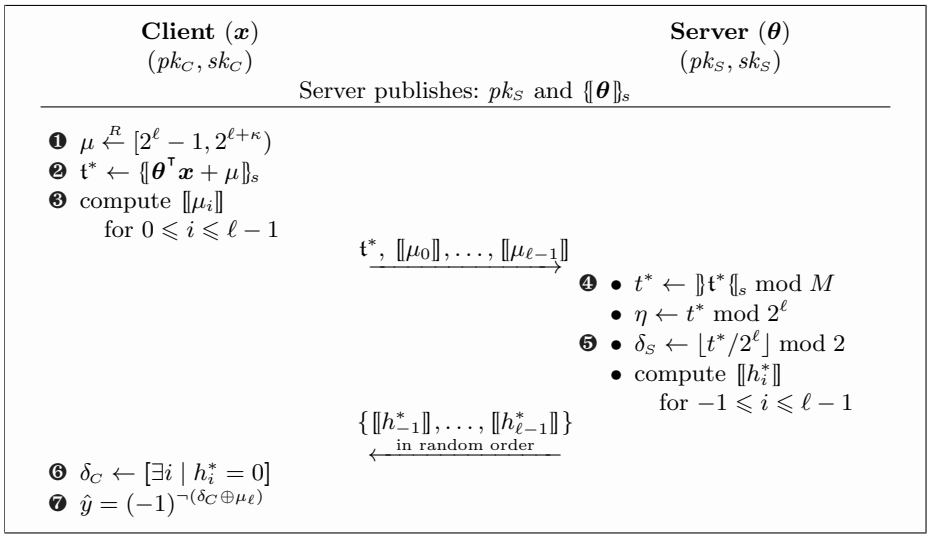

Fig. 4: Privacy-preserving SVM classification. Note that some data is encrypted using the client's public key pk_C, while other, is encrypted using the server's public key pk_S. They are noted $[\![\cdot]\!]$ and $\{\!|\cdot|\!\}_s$ respectively.

Correctness Remember that, by construction, $\boldsymbol{\theta}^\intercal \boldsymbol{x} \in [-B, B]$ with $B = 2^\ell - 1$, that $\mu \in [2^\ell - 1, 2^{\ell+\kappa})$, and by definition that $t^* := |\!|\mathfrak{t}^*|\!|_s \bmod M$ with $\mathfrak{t}^* = \{\!|\boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu|\!\}_s$. Hence, in Step 4, the server gets $t^* = \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu \bmod M = \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$ (over $\mathbb{Z}$) since $0 \leqslant \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu \leqslant 2^\ell - 1 + 2^{\ell+\kappa} - 1 < M$. Let $\delta := \delta_C \oplus \delta_S = [\underline{\mu} \leqslant \eta]$ (with $\underline{\mu} := \mu \bmod 2^\ell$ and $\eta := t^* \bmod 2^\ell$) denote the result of the private comparison in Steps 5 and 6 with the DGK+ protocol.

Either of those two conditions holds true

$$\begin{cases} 0 \leqslant \boldsymbol{\theta}^\intercal \boldsymbol{x} < 2^\ell & \iff 1 \leqslant \frac{\boldsymbol{\theta}^\intercal \boldsymbol{x} + 2^\ell}{2^\ell} < 2 \iff \left\lfloor \frac{\boldsymbol{\theta}^\intercal \boldsymbol{x} + 2^\ell}{2^\ell} \right\rfloor = 1 \\ -2^\ell < \boldsymbol{\theta}^\intercal \boldsymbol{x} < 0 & \iff 0 < \frac{\boldsymbol{\theta}^\intercal \boldsymbol{x} + 2^\ell}{2^\ell} < 1 \iff \left\lfloor \frac{\boldsymbol{\theta}^\intercal \boldsymbol{x} + 2^\ell}{2^\ell} \right\rfloor = 0 \end{cases},$$

and so, since $t^* = \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$, $[\boldsymbol{\theta}^\intercal \boldsymbol{x} \geqslant 0] \in \{0, 1\}$, and $\delta_S = \lfloor t^*/2^\ell \rfloor \bmod 2$:

$$\begin{aligned} [\boldsymbol{\theta}^\intercal \boldsymbol{x} \geqslant 0] &= \left\lfloor \tfrac{\boldsymbol{\theta}^\intercal \boldsymbol{x} + 2^\ell}{2^\ell} \right\rfloor = \left\lfloor \tfrac{t^* - \mu}{2^\ell} \right\rfloor + 1 = \left\lfloor \tfrac{t^*}{2^\ell} \right\rfloor - \left\lfloor \tfrac{\mu}{2^\ell} \right\rfloor + \left\lfloor \tfrac{\eta - \mu}{2^\ell} \right\rfloor + 1 \\ &= \left\lfloor \tfrac{t^*}{2^\ell} \right\rfloor - \left\lfloor \tfrac{\mu}{2^\ell} \right\rfloor + \delta = \left(\left\lfloor \tfrac{t^*}{2^\ell} \right\rfloor - \left\lfloor \tfrac{\mu}{2^\ell} \right\rfloor + \delta \right) \bmod 2 \\ &= \left(\left\lfloor \tfrac{\mu}{2^\ell} \right\rfloor + \delta_C \right) \bmod 2 = \mu_\ell \oplus \delta_C \ . \end{aligned}$$

Now, noting $\mathrm{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x}) = (-1)^{\neg[\boldsymbol{\theta}^\intercal \boldsymbol{x} \geqslant 0]}$, we get the desired result.

Security The security of the protocol of Fig. 4 follows from the fact that the inner product $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ is statistically masked by the random value μ. Security parameter κ guarantees that the probability of an information leak due to a carry is negligible. The security also depends on the security of the DGK+ comparison protocol, which is provably secure (cf. Remark 2).

A Heuristic Protocol The previous protocol, thanks to the use of the DGK+ algorithm offers provable security guarantees but incurs the exchange of $2(\ell+1)$ ciphertexts. Here we aim to reduce the number of ciphertexts and introduce a new heuristic protocol. This protocol requires the introduction of a signed factor λ, such that $|\lambda| > |\mu|$, and we now use both μ and λ to mask the model. To ensure that $\lambda\boldsymbol{\theta}^\intercal\boldsymbol{x}+\mu$ remains within the message space, we pick λ in $\mathcal{B}$ where $\mathcal{B} \coloneqq \left[-\left\lceil\frac{\lceil M/2\rceil}{B+1}\right\rceil, \left\lfloor\frac{\lceil M/2\rceil}{B+1}\right\rfloor\right]$. Furthermore, to ensure the effectiveness of the masking, $\mathcal{B}$ should be sufficiently large; namely, $\#\mathcal{B} > 2^\kappa$ for a security parameter κ, hence $M > 2^\ell(2^\kappa-1)$.

The protocol, which is illustrated in Fig. 5, runs as follows:

1. The client encrypts its input data $\boldsymbol{x}$ using its public key, and sends its key and the encrypted data to the server.
2. The server draws at random a signed scaling factor $\lambda \in \mathcal{B}$, $\lambda \neq 0$, and an offset factor $\mu \in \mathcal{B}$ such that $|\mu| < |\lambda|$ and $\text{sign}(\mu) = \text{sign}(\lambda)$. The server then defines the bit δ_S such that $\text{sign}(\lambda) = (-1)^{\delta_S}$ and computes an encryption $\mathfrak{t}^*$ of the shifted and scaled inner product $t^* = (-1)^{\delta_S}\cdot(\lambda\boldsymbol{\theta}^\intercal\boldsymbol{x}+\mu)$ and sends $\mathfrak{t}^*$ to the client.
3. In the final step, the client decrypts $\mathfrak{t}^*$ using its private key, recovers t^* as a signed integer of $\mathcal{M}$, and deduces the class of the input data as $\hat{y} = \text{sign}(t^*)$.

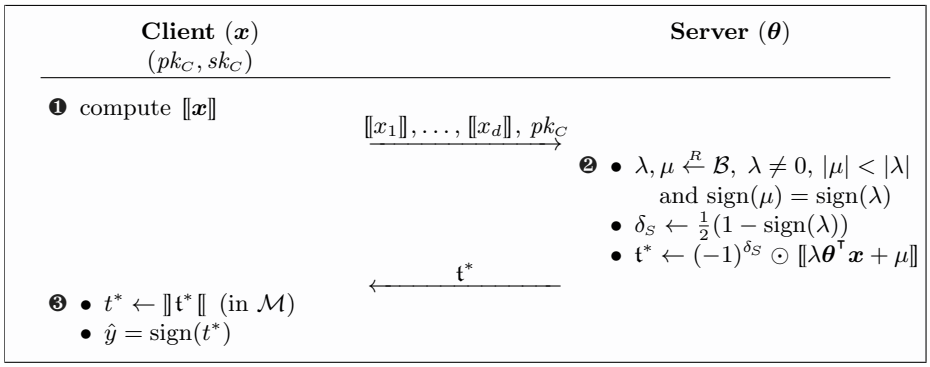

Fig. 5: Heuristic protocol for privacy-preserving SVM classification.

Correctness The constraint $|\mu| < |\lambda|$ with $\lambda \neq 0$ ensures that $\hat{y} = \text{sign}(\boldsymbol{\theta}^\intercal\boldsymbol{x})$. Indeed, as $(-1)^{\delta_S} = \text{sign}(\lambda) = \text{sign}(\mu)$, we have $t^* = (-1)^{\delta_S}(\lambda\boldsymbol{\theta}^\intercal\boldsymbol{x}+\mu) = |\lambda|\boldsymbol{\theta}^\intercal\boldsymbol{x}+|\mu| = |\lambda|(\boldsymbol{\theta}^\intercal\boldsymbol{x}+\epsilon)$ with $\epsilon \coloneqq |\mu|/|\lambda|$. Hence, whenever $\boldsymbol{\theta}^\intercal\boldsymbol{x} \neq 0$, we get $\hat{y} = \text{sign}(t^*) = \text{sign}(\boldsymbol{\theta}^\intercal\boldsymbol{x}+\epsilon) = \text{sign}(\boldsymbol{\theta}^\intercal\boldsymbol{x})$ since $|\boldsymbol{\theta}^\intercal\boldsymbol{x}| \geqslant 1$ and $|\epsilon| = |\mu|/|\lambda| < 1$. If $\boldsymbol{\theta}^\intercal\boldsymbol{x} = 0$ then $\hat{y} = \text{sign}(\epsilon) = 1$.

Security We stress that the private comparison protocol we use in Fig. 5 does not come with formal security guarantees. In particular, the client learns the

value of $t^* = \lambda \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$ with $\lambda, \mu \in \mathcal{B}$ and $|\mu| < |\lambda|$. Some information on $t \coloneqq \boldsymbol{\theta}^\intercal \boldsymbol{x}$ may be leaking from t^* and, in turn, on $\boldsymbol{\theta}$ since $\boldsymbol{x}$ is known to the client. The reason resides in the constraint $|\mu| < |\lambda|$. So, from $t^* = \lambda \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$, we deduce $\log|t^*| \leqslant \log|\lambda| \cdot \log(|t| + 1)$. For example, when t has two possible very different "types" of values (say, very large and very small), the quantity $\log|t^*|$ can be enough to discriminate with non-negligible probability the type of t. This may possibly leak information on $\boldsymbol{\theta}$. That does not mean that the protocol is necessarily insecure but it should be used with care.

4 Application to Neural Networks

Typical *feed-forward neural networks* are represented as large graphs. Each node on the graph is often called a *unit*, and these units are organised into layers. At the very bottom is the input layer with a unit for each of the coordinates $x_j^{(0)}$ of the input vector $\boldsymbol{x}^{(0)} \coloneqq \boldsymbol{x} \in \mathcal{X}$. Then various computations are done in a bottom-to-top pass and the output $\hat{y} \in \mathcal{Y}$ comes out all the way at the very top of the graph. Between the input and output layers, a number of *hidden layers* are evaluated. We index the layers with a superscript (l), where $l = 0$ for the input layer and $1 \leqslant l < L$ for the hidden layers. Layer L corresponds to the output. Each unit of each layer has directed connections to the units of the layer below; see Fig. 6.

We keep the convention $x_0^{(l)} \coloneqq 1$ for all layers. If we note $\boldsymbol{\theta_j}^{(l)}$ the vector of weight coefficients $\theta_{j,k}^{(l)}$, $0 \leqslant k \leqslant d_{l-1}$, where d_l is the number of units in layer l, then $x_j^{(l)}$ can be expressed as:

$$x_j^{(l)} = g_j^{(l)}\left(\left(\boldsymbol{\theta_j}^{(l)}\right)^\intercal \boldsymbol{x}^{(l-1)}\right), \quad 1 \leqslant j \leqslant d_l \ . \tag{2}$$

Functions $g_j^{(l)}$ are non-linear functions such as the sign function or the Rectified Linear Unit (ReLU) function (see Section 4.2). Those functions are known as *activation functions*. The weight coefficients characterise the model and are

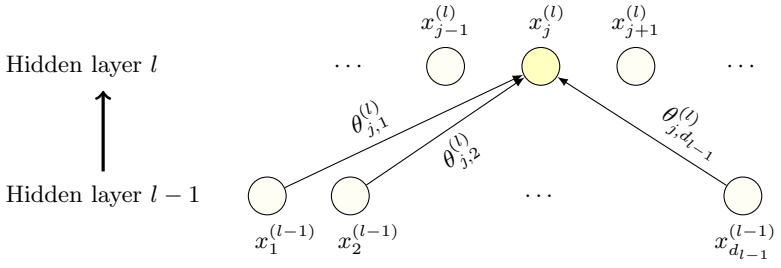

Fig. 6: Relationship between a hidden unit in layer l and the hidden units of layer $l-1$ in a simple feed-forward neural network.

known only to the owner of the model. Each hidden layer depends on the layer below, and ultimately on the input data $\boldsymbol{x}^{(0)}$, known solely to the client.

Generic Solution A generic solution can easily be devised from Eq. (2): for each inner product computation, and therefore for each unit of each hidden layer, the server computes the encrypted inner product and the client computes the output of the activation function in the clear. In more detail, the evaluation of a neural network can go as follows.

0. The client starts by encrypting its input data and sends it to the server.
1. Then, as illustrated in Fig. 7, for each hidden layer l, $1 \leqslant l < L$:
   (a) The server computes d_l encrypted inner products t_j corresponding to each unit j of the layer and sends those to the client.
   (b) The client decrypts the inner products, applies the required activation function $g_j^{(l)}$, re-encrypts, and sends back d_l encrypted values.
2. During the last round ($l = L$), the client simply decrypts the t_j values and applies the corresponding activation function $g_j^{(L)}$ to each unit j of the output layer. This is the required result.

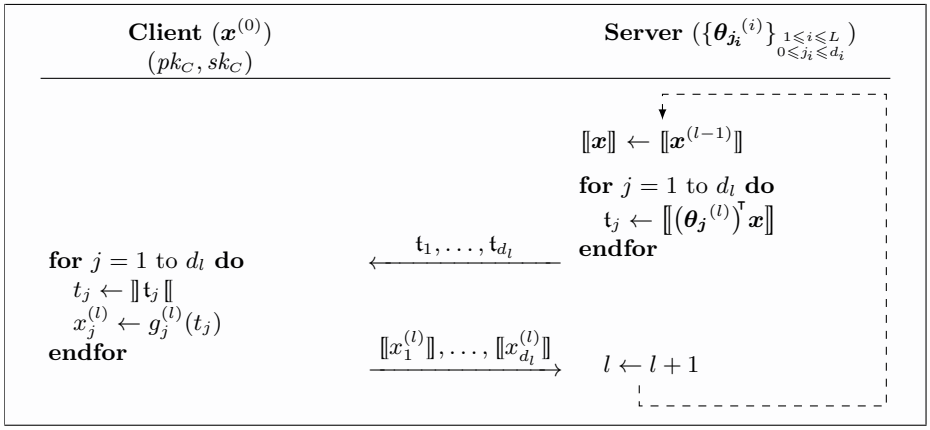

Fig. 7: Generic solution for privacy-preserving evaluation of feed-forward neural networks. Evaluation of hidden layer l.

For each hidden layer l, exactly two messages (each comprising d_l encrypted values) are exchanged. The input and output layers only involve one exchange; from the client to the server for the input layer and from the server back to the client for the output layer.

In the following two sections, we improve this generic solution for two popular activation functions: the sign and the ReLU functions. In the new proposed implementations, everything is kept encrypted—from start to end. The raw signals are hidden from the client's view in all intermediate computations.

4.1 Sign Activation

Binarised neural networks implement the sign function as activation function. This is very advantageous from a hardware perspective [14].

Section 3.2 describes two protocols for the client to get the sign of $\boldsymbol{\theta}^\intercal \boldsymbol{x}$. In order to use them for binarised neural networks in a setting similar to the generic solution, the server needs to get an encryption of $\mathrm{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x})$ for each computing unit j in layer l under the client's key from $[\![\boldsymbol{x}]\!]$, where $[\![\boldsymbol{x}]\!] \coloneqq [\![\boldsymbol{x}^{(l-1)}]\!]$ is the encrypted output of layer $l-1$ and $\boldsymbol{\theta} \coloneqq \boldsymbol{\theta_j}^{(l)}$ is the parameter vector for unit j in layer l.

We start with the core protocol of Fig. 4. It runs in dual mode and therefore uses the server's encryption. Exchanging the roles of the client and the server almost gives rise to the sought-after protocol. The sole extra change is to ensure that the server gets the classification result encrypted. This can be achieved by masking the value of δ_C with a random bit b and sending an encryption of $(-1)^b$. The resulting protocol is depicted in Fig. 8.

In the heuristic protocol (cf. Fig. 5), the server already gets an encryption of $[\![\boldsymbol{x}]\!]$ as an input. It however fixes the sign of t^* to that of $\boldsymbol{\theta}^\intercal \boldsymbol{x}$. If now the

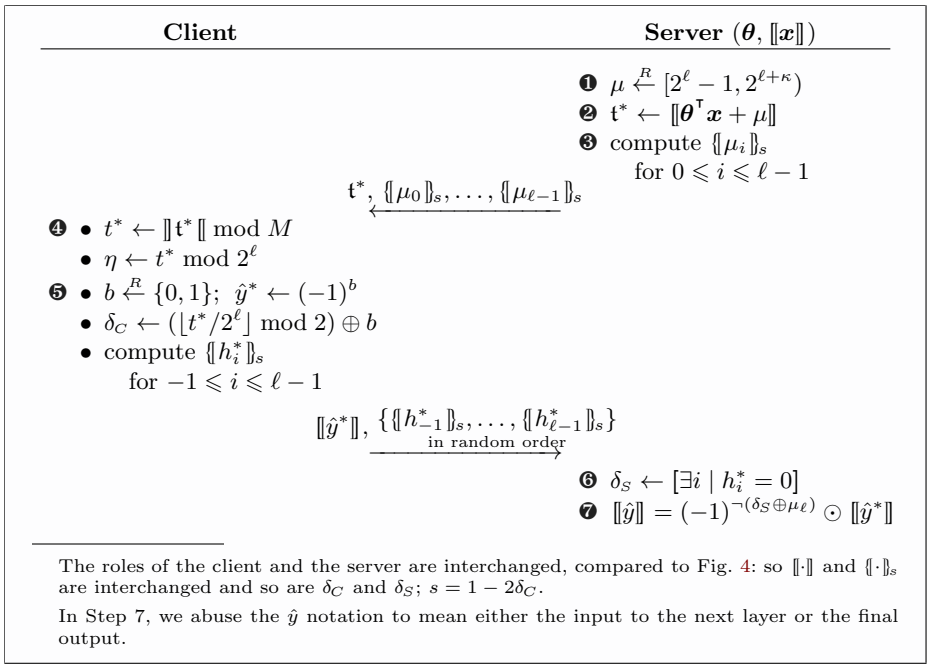

The roles of the client and the server are interchanged, compared to Fig. 4: so $[\![\cdot]\!]$ and $\{\!|\cdot|\!\}_s$ are interchanged and so are δ_C and δ_S; $s = 1 - 2\delta_C$.

In Step 7, we abuse the $\hat{y}$ notation to mean either the input to the next layer or the final output.

Fig. 8: Privacy-preserving binary classification with inputs and outputs encrypted under the client's public key. This serves as a building block for the evaluation over encrypted data of the sign activation function in a neural network and shows the computations and message exchanges for one unit in one hidden layer.

server flips it in a probabilistic manner, the output class (i.e., $\mathrm{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x})$) will be hidden from the client's view. We detail below the modifications to be brought to the heuristic protocol to accommodate the new setting:

- In Step 2 of Fig. 5, the server keeps private the value of δ_S by replacing the definition of $\mathfrak{t}^*$ with $\mathfrak{t}^* = [\![\lambda \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu]\!]$.
- In Step 3 of Fig. 5, the client then obtains $\hat{y}^* \coloneqq \mathrm{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x}) \cdot (-1)^{\delta_S}$ and returns its encryption $[\![\hat{y}^*]\!]$ to the server.
- The server obtains $[\![\hat{y}]\!]$ as $[\![\hat{y}]\!] = (-1)^{\delta_S} \odot [\![\hat{y}^*]\!]$.

If $\boldsymbol{\theta} \coloneqq \boldsymbol{\theta_j}^{(l)}$ and $[\![\boldsymbol{x}]\!] \coloneqq [\![\boldsymbol{x}^{(l)}]\!]$ then the outcome of the protocol of Fig. 8 or of the modified heuristic protocol is $[\![\hat{y}]\!] = [\![x_j^{(l)}]\!]$. Of course, this can be done in parallel for all the d_l units of layer l (i.e., for $1 \leqslant j \leqslant d_l$; see Eq. (2)), yielding $[\![\boldsymbol{x}^{(l)}]\!] = ([\![1]\!], [\![x_1^{(l)}]\!], \ldots, [\![x_{d_l}^{(l)}]\!])$. This means that just one round of communication between the server and the client suffices per hidden layer.

4.2 ReLU Activation

A widely used activation function is the ReLU function. It allows a network to easily obtain sparse representations and features cheaper computations as there is no need for computing the exponential function [10].

The ReLU function can be expressed from the sign function as

$$\mathrm{ReLU}(t) = \tfrac{1}{2}(1 + \mathrm{sign}(t)) \cdot t \ . \tag{3}$$

Back to our setting, the problem is for the server to obtain $[\![\mathrm{ReLU}(t)]\!]$ from $[\![t]\!]$, where $t = \boldsymbol{\theta}^\intercal \boldsymbol{x}$ with $\boldsymbol{x} \coloneqq \boldsymbol{x}^{(l-1)}$ and $\boldsymbol{\theta} \coloneqq \boldsymbol{\theta_j}^{(l)}$, in just one round of communication per hidden layer. We saw in the previous section how to do it for the sign function. The ReLU function is more complex to apprehend. If we use Equation (3), the difficulty is to let the server evaluate a *product* over encrypted data. To get around that, the server super-encrypts $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!]$, gets $\{\!\{ [\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!] \}\!\}_s$, and sends it the client. According to its secret share the client sends back the pair $(\{\!\{ [\![0]\!] \}\!\}_s, \{\!\{ [\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!] \}\!\}_s)$ or $(\{\!\{ [\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!] \}\!\}_s, \{\!\{ [\![0]\!] \}\!\}_s)$. The server then uses its secret share to select the correct item in the received pair, decrypts it, and obtains $[\![\mathrm{ReLU}(\boldsymbol{\theta}^\intercal \boldsymbol{x})]\!]$. For this to work, it is important that the client re-randomises $\{\!\{ [\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!] \}\!\}_s$ as otherwise the server could distinguish it from $\{\!\{ [\![0]\!] \}\!\}_s$.

Actually, a simple one-time pad suffices to implement the above solution. To do so, the server chooses a random mask $\mu \in \mathcal{M}$ and "super-encrypts" $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!]$ as $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu]\!]$. The client re-randomises it as $\mathfrak{t}^{**} \coloneqq [\![\boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu]\!] \boxplus [\![0]\!]$, computes $\mathfrak{o} \coloneqq [\![0]\!]$, and returns the pair $(\mathfrak{o}, \mathfrak{t}^{**})$ or $(\mathfrak{t}^{**}, \mathfrak{o})$, depending on its secret share. The server uses its secret share to select the correct item and "decrypts" it. If the server (obliviously) picked $\mathfrak{o}$, it already has the result in the right form; i.e., $[\![0]\!]$. Otherwise the server has to remove the mask μ so as to get $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!] \leftarrow \mathfrak{t}^{**} \boxminus [\![\mu]\!]$. In order to allow the server to (obliviously) remove or not the mask, the client also sends an encryption of the pair index; e.g., 0 for the pair $(\mathfrak{o}, \mathfrak{t}^{**})$ and 1 for the pair $(\mathfrak{t}^{**}, \mathfrak{o})$.

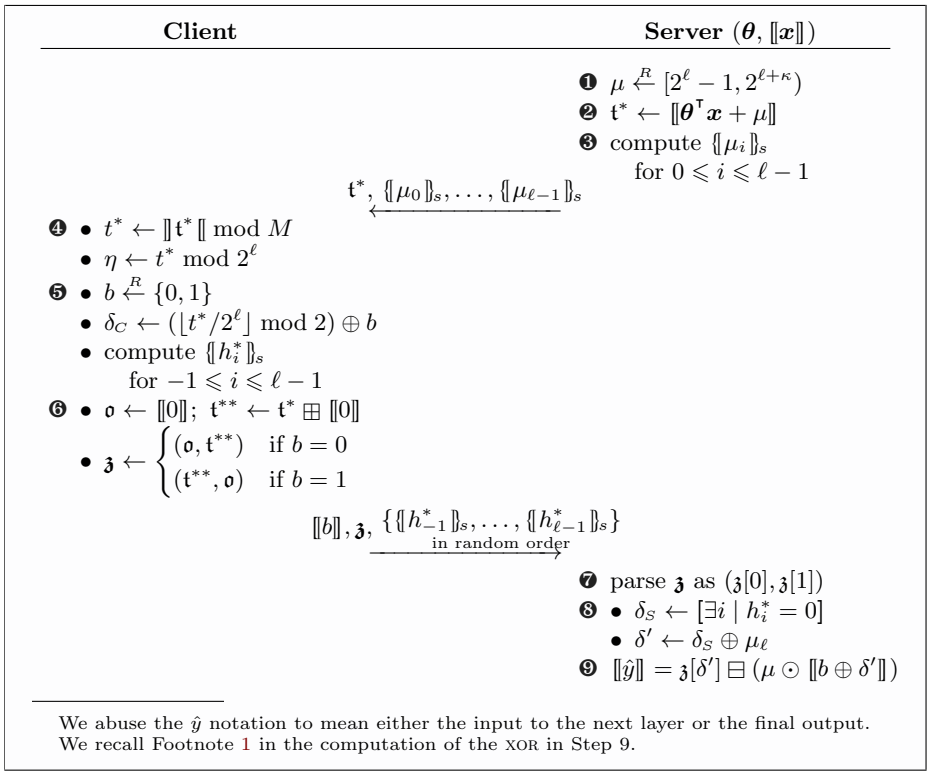

We abuse the $\hat{y}$ notation to mean either the input to the next layer or the final output. We recall Footnote 1 in the computation of the XOR in Step 9.

Fig. 9: Privacy-preserving ReLU evaluation with inputs and outputs encrypted under the client's public key. The first five steps are as in Fig. 8. This building block is directed to neural networks using the ReLU activation and shows the computations and messages exchanges for one unit in one hidden layer.

Figure 9 details an implementation of this with the DGK+ comparison protocol. Note that to save on bandwidth the same mask μ is used for the comparison protocol and to "super-encrypt" $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!]$. The heuristic protocol can be adapted in a similar way.

Remark 3. It is interesting to note that the new protocols readily extend to any piece-wise linear function, such as the clip function $\text{clip}(t) = \max(0, \min(1, \frac{t+1}{2}))$ (a.k.a. hard-sigmoid function).

5 Conclusion

In this work, we presented several protocols for privacy-preserving regression and classification. Those protocols only require additively homomorphic encryption and limit interactions to a mere request and response. They are secure against

semi-honest adversaries. They can be used as-is in generalised linear models (including logistic regression and SVM classification) or applied to other machine-learning algorithms. As an illustration, we showed how they nicely adapt to binarised neural networks or to feed-forward neural networks with the ReLU activation function.

References

1. Abu-Mostafa, Y.S., Magdon-Ismail, M., Lin, H.T.: Learning From Data: A Short Course. AMLbook.com (2012)
2. Agrawal, R., Srikant, R.: Privacy-preserving data mining. ACM Sigmod Record **29**(2), 439–450 (2000)
3. Barni, M., Orlandi, C., Piva, A.: A privacy-preserving protocol for neural-network-based computation. In: MM&Sec 2006. pp. 146–151. ACM (2006)
4. Bos, J.W., Lauter, K., Naehrig, M.: Private predictive analysis on encrypted medical data. J. Biomed. Inf. **50**, 234–243 (2014)
5. Bost, R., Popa, R.A., Tu, S., Goldwasser, S.: Machine learning classification over encrypted data. In: NDSS 2015. The Internet Society (2015)
6. Damgård, I., Geisler, M., Krøigaard, M.: Homomorphic encryption and secure comparison. Int. J. Appl. Cryptogr. **1**(1), 22–31 (2008)
7. Damgård, I., Geisler, M., Krøigaard, M.: A correction to 'efficient and secure comparison for on-line auctions'. Int. J. Appl. Cryptogr. **1**(4), 323–324 (2009)
8. Dwork, C., Feldman, V.: Privacy-preserving prediction. In: COLT 2018. PMLR, vol. 75, pp. 1693–1702. PMLR (2018)
9. Erkin, Z., Franz, M., Guajardo, J., Katzenbeisser, S., Lagendijk, I., Toft, T.: Privacy-preserving face recognition. In: PETS 2009. LNCS, vol. 5672, pp. 235–253. Springer (2009)
10. Glorot, X., Bordes, A., Bengjio, Y.: Deep sparse rectifier neural networks. In: AISTAT 2011. PMLR, vol. 15, pp. 315–323. PMLR (2011)
11. Goethals, B., Laur, S., Lipmaa, H., Mielikäinen, T.: On private scalar product computation for privacy-preserving data mining. In: ICISC 2004. LNCS, vol. 3506, pp. 104–102. Springer (2004)
12. Goldwasser, S., Micali, S.: Probabilistic encryption. J. Comput. Syst. Sci. **28**(2), 270–299 (1984)
13. Hastie, T., Tibshirani, R., Friedman, J.: The Elements of Statistical Learning. Springer Series in Statistics, Springer, 2nd edn. (2009)
14. Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks. In: NISP 2016. pp. 4107–4115 (Curran Associates, Inc)
15. Joye, M., Salehi, F.: Private yet efficient decision tree evaluation. In: DBSEC 2018. LNCS, vol. 10980, pp. 243–259. Springer (2018)
16. Kim, M., Song, Y., Wang, S., Xia, Y., Jiang, X.: Secure logistic regression based on homomorphic encryption: Design and evaluation. JMIR Med. Inform. **6**(2), e19 (2018)
17. Lindell, Y., Pinkas, B.: Privacy preserving data mining. In: CRYPTO 2000. LNCS, vol. 1880, pp. 36–54. Springer (2000)
18. Mohassel, P., Zhang, Y.: SecureML: A system for scalable privacy-preserving machine learning. In: IEEE S&P 2017. pp. 19–38. IEEE Computer Society (2017)
19. Tramèr, F., Zhang, F., Juels, A., Reiter, M.K., Ristenpart, T.: Stealing machine learning models via prediction APIs. In: USENIX Security 2016. pp. 601–618. USENIX Association (2016)

20. Veugen, T.: Improving the DGK comparison protocol. In: WIFS 2012. pp. 49–54. IEEE (2012)
21. Zhang, J., Wang, X., Yiu, S.M., Jiang, Z.L., Li, J.: Secure dot product of outsourced encrypted vectors and its application to SVM. In: SCC@AsiaCCS 2017. pp. 75–82. ACM (2017)

A More Private Protocols

A.1 Linear Regression

As seen in Section 2.1, linear regression produces estimates using the identity map for g: $\hat{y} = \boldsymbol{\theta}^\intercal \boldsymbol{x}$. Since $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ is linear, given an encryption $[\![\boldsymbol{x}]\!]$ of $\boldsymbol{x}$, the value of $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!]$ can be homomorphically evaluated, in a provably secure way [11].

Therefore, the client encrypts its feature vector $\boldsymbol{x}$ under its public key and sends $[\![\boldsymbol{x}]\!]$ to the server. Using $\boldsymbol{\theta}$, the server then computes $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!]$ and returns it the client. Finally, the client decrypts $[\![\boldsymbol{\theta}^\intercal \boldsymbol{x}]\!]$ and gets the output $\hat{y} = \boldsymbol{\theta}^\intercal \boldsymbol{x}$. This is straightforward and only requires one round of communication.

A.2 SVM Classification

A Naïve Protocol A client holding a private feature vector $\boldsymbol{x}$ wishes to evaluate $\operatorname{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x})$ where $\boldsymbol{\theta}$ parametrises an SVM classification model. In the primal approach, the client can encrypt $\boldsymbol{x}$ and send $[\![\boldsymbol{x}]\!]$ to the server. Next, the server computes $[\![\eta]\!] = [\![\boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu]\!]$ for some random mask μ and sends $[\![\eta]\!]$ to the client. The client decrypts $[\![\eta]\!]$ and recovers η. Finally, the client and the server engage in a private comparison protocol with respective inputs η and μ, and the client deduces the sign of $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ from the resulting comparison bit $[\mu \leqslant \eta]$.

There are two issues. If we use the DGK+ protocol for the private comparison, at least one extra exchange from the server to the client is needed for the client to get $[\mu \leqslant \eta]$. This can be fixed by considering the dual approach. A second, more problematic, issue is that the decryption of $[\![\eta]\!] \coloneqq [\![\boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu]\!]$ yields η as an element of $\mathcal{M} \cong \mathbb{Z}/M\mathbb{Z}$, which is not necessarily equivalent to the *integer* $\boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu$. Note that if the inner product $\boldsymbol{\theta}^\intercal \boldsymbol{x}$ can take any value in $\mathcal{M}$, selecting a smaller value for $\mu \in \mathcal{M}$ to prevent the modular reduction does not solve the issue because the value of η may then leak information on $\boldsymbol{\theta}^\intercal \boldsymbol{x}$.

A Heuristic Protocol (Dual Approach) The bandwidth usage with the heuristic comparison protocol (cf. Fig. 5) could be even reduced to one ciphertext and a single bit with the dual approach. From the published encrypted model $\{\!|\boldsymbol{\theta}|\!\}_s$, the client could homomorphically compute and send to the server $\mathfrak{t}^* = \{\!|\lambda \boldsymbol{\theta}^\intercal \boldsymbol{x} + \mu|\!\}_s$ for random $\lambda, \mu \in \mathcal{B}$ with $|\mu| < |\lambda|$. The server would then decrypt $\mathfrak{t}^*$, obtain t^*, compute $\delta_S = \frac{1}{2}(1 - \operatorname{sign}(t^*))$, and return δ_S to the client. Analogously to the primal approach, the output class $\hat{y} = \operatorname{sign}(\boldsymbol{\theta}^\intercal \boldsymbol{x})$ is obtained by the client as $\hat{y} = (-1)^{\delta_S} \cdot \operatorname{sign}(\lambda)$. However, and contrarily to the primal approach, the potential information leakage resulting from t^*—in this case on $\boldsymbol{x}$—is now on the server's side, which is in contradiction with our Requirement 1 (input confidentiality). We do not further discuss this variant.