

Secure PAC Bayesian Regression via Real Shamir Secret Sharing

Jaron Skovsted Gundersen, Bulut Kuskonmaz, Rafael Wisniewski

Abstract—A common approach of system identification and machine learning is to generate a model by using training data to predict the test data instances as accurately as possible. Nonetheless, concerns about data privacy are increasingly raised, but not always addressed. We present a secure protocol for learning a linear model relying on a recently described technique called real number secret sharing. We take as our starting point the PAC Bayesian bounds and deduce a closed form for the model parameters which depends on the data and the prior from the PAC Bayesian bounds. To obtain the model parameters one needs to solve a linear system. However, we consider the situation where several parties hold different data instances and they are not willing to give up the privacy of the data. Hence, we suggest to use real number secret sharing and multiparty computation to share the data and solve the linear regression with a secure distributed Gaussian elimination protocol such that privacy of the data is preserved. The benefit of using secret sharing directly on real numbers is reflected in the simplicity of the protocols and the number of rounds needed. However, this comes with the drawback that a share might leak a small amount of information, but in our analysis we argue that the leakage is small.

I. INTRODUCTION

The main purpose of system identification and machine learning is to find a model based on data such that the model generalizes to instances outside the data set. The more data we possess the better is the model. However, the data may be sensitive and hence we are not allowed to share the data with others. In this paper, we consider the case where several parties hold a data set and each party wants to learn a model trained on the data in all the data sets. We strive to learn the model without giving up the privacy of the data. One way to determine a model is to rely on the PAC Bayesian framework, which can be used if some prior knowledge of the model is available. There are a lot of PAC Bayesian bounds in the literature which assess the generalization error with a high probability, see for instance [2, 5, 16, 18]. We use the PAC Bayesian framework with a normally distributed prior to obtain a linear model estimation. In this setup, we get a closed form solution for the estimation.

For the purpose of privacy, techniques from secure multiparty computation (MPC) are useful. In [19], they apply homomorphic encryption and Yao garbled circuits to achieve the secure ridge regression as a secure machine learning algorithm. In [8] MPC is used to securely compute linear regression without prior and where the data is split between two parties. Another study uses encrypted data for their classification phases as they call it privacy-preserving classification [4]. Originally, and in these papers, MPC is carried out over a finite field, but recent research considers MPC over

real numbers [22] and we take this approach. However, this is a fairly new approach and in [22] the privacy is only shown for a single operation. In this paper we want to argue that the privacy is retained even when several secure computations are carried out sequentially.

With the assumption that data is split between different parties, we show how one can use real number secret sharing and MPC techniques to compute the model without violating privacy. To argue privacy, we consider the notion of conditional mutual information $I(X; V|Y)$. In our setup Y is everything that an adversary will learn in an ideal world (the data of the corrupted parties and the model), X is what we want to hide (the data of the honest parties), and V is the additional information an adversary gets by using the protocol. Thus, $I(X; V|Y)$ describes how much more information the adversary learns about the honest parties data by using the protocol compared to an ideal version where only the model is given to them. Ideally this quantity should be zero which is equivalent to that X and V are conditionally independent given Y . Since $I(X; V|Y)$ is hard to determine in our case, due to the high dimensions of X, V, Y , we will use a conditionally independent test to argue privacy. Conditionally independent tests are a line of research of its own. More recently, kernel based and neural network based tests have been developed which is argued to be applicable in the high dimensional cases [23, 3].

II. PRELIMINARIES

A. Notation

Let $\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \dots \cup \mathcal{D}_n$ be a split data set and assume that party i holds

$$\mathcal{D}_i = \{(\mathbf{x}_{i1}, y_{i1}), (\mathbf{x}_{i2}, y_{i2}), \dots, (\mathbf{x}_{in_i}, y_{in_i})\}.$$

Furthermore, we denote by $N = \sum_{i=1}^n n_i = |\mathcal{D}|$ and assume that $(\mathbf{x}_{ij}, y_{ij}) \in \mathbb{R}^d \times \mathbb{R}$ where d is the dimension of data.

Assume that a model f is given. For a data point $(\mathbf{x}, y) \in \mathcal{D}$, the prediction of \mathbf{x} is $f(\mathbf{x})$ and the squared loss is denoted by $\ell(f(\mathbf{x}), y) = (f(\mathbf{x}) - y)^2$. The empirical loss is

$$\mathcal{L}_N(f) = \frac{1}{N} \sum_{i=1}^n \sum_{j=1}^{n_i} \ell(f(\mathbf{x}_{ij}), y_{ij}), \quad (1)$$

and the generalization error is the expected value of the loss function, i.e., $\mathcal{L}(f) = E_{(\mathbf{x}, y)}[\ell(f(\mathbf{x}), y)]$, where the expectation is taken over the joint distribution of \mathbf{x} and y .

B. Multiparty computation

MPC deals with the situation where there are n parties, and each has some inputs to a function. The parties would

like to learn the output of the function without revealing their inputs to the other parties. An MPC protocol should make it possible for the parties to obtain this even in the presence of an adversary corrupting some of the parties. In this paper, we consider a passive corrupt adversary, meaning that the adversary can see everything the corrupted parties are able to see but they cannot deviate from the protocol description. To prove that a protocol is secure we usually compare what harm an adversary can do in an execution of the protocol to what harm an adversary can do in an idealized version. Since, we are considering a passive adversary the only “harm” is that the adversary learns the information it was not supposed to learn. However, this can also be formulated through mutual information¹. Let X and Y be random variables then the mutual information can be defined

$$I(X; Y) = h(X) - h(X|Y)$$

where $h(X)$ is the entropy of X and $h(X|Y)$ is the conditional entropy.² With this definition we use the same definition as in [17] to define privacy of a protocol.

1 Definition (Perfect Privacy):

Let *View* be anything the adversary sees through a protocol computing $f(x_1, \dots, x_n) = y$ and let \mathcal{A} and \mathcal{H} be the indices of the inputs held by corrupted and honest nodes. Then we say that the protocol obtains perfect privacy if

$$I(\{x_i\}_{i \in \mathcal{H}}; \text{View}) = I(\{x_i\}_{i \in \mathcal{H}}; \{y\} \cup \{x_j\}_{j \in \mathcal{A}}).$$

When considering continuous random variables, $I(\{x_i\}_{i \in \mathcal{H}}; \text{View})$ might be very difficult to compute since we are rarely able to determine the distribution on *View* and it will often have a very high dimension.³

Since the mutual information is difficult to compute in our setup, for instance because of the random variables have high dimensions, we take another approach. Note that $\text{View} = V \cup \{y\} \cup \{x_j\}_{j \in \mathcal{A}}$ where V includes all the randomness chosen by the adversary and all the messages the adversary receives. Since we do not obtain perfect privacy we want to show that V does not reveal much information about $\{x_i\}_{i \in \mathcal{H}}$. That is we want

$$\begin{aligned} I(\{x_i\}_{i \in \mathcal{H}}; \text{View}) - I(\{x_i\}_{i \in \mathcal{H}}; \{y\} \cup \{x_j\}_{j \in \mathcal{A}}) \\ = I(\{x_i\}_{i \in \mathcal{H}}; V | \{y\} \cup \{x_j\}_{j \in \mathcal{A}}). \end{aligned} \quad (2)$$

to be small. If $\{x_i\}_{i \in \mathcal{H}}$ and V are conditionally independent conditioned on $\{y\} \cup \{x_j\}_{j \in \mathcal{A}}$ this conditional mutual information is zero. Thus in our privacy analysis we apply a conditional independence test to show that if we apply our distributed algorithm on a given data set with large enough security parameters we cannot reject the null hypothesis

¹For more information about mutual information we refer to [9].

²Mutual information can also equivalently be defined through the Kullback-Leibler divergence where the joint distribution of X and Y is compared to the product of the marginals.

³This is for instance mentioned in [15] where alternatives to the mutual information is suggested by using another measure than the KL-divergence. However, these suggestions have also computational problems when the dimension increases.

that $\{x_i\}_{i \in \mathcal{H}}$ and V are conditionally independent given $\{y\} \cup \{x_j\}_{j \in \mathcal{A}}$. And hence $I(\{x_i\}_{i \in \mathcal{H}}; V | \{y\} \cup \{x_j\}_{j \in \mathcal{A}})$ must be close to zero. We apply the test from [3].

Many MPC protocols are secret sharing-based meaning that each party starts by sending a share of its input to all the other parties. The share itself does not reveal information and in fact we say that the secret sharing scheme has privacy threshold t , if t shares do not reveal anything about the secret. However, if more than t shares are collected we should be able to reconstruct the secret. MPC can use properties for certain secret sharing schemes allowing us to do computations on the shares. From these computations, we end up with the parties having a share for the output. By broadcasting these shares, they can by the reconstruction property in the secret sharing scheme compute the output. For a comprehensive introduction to secret sharing and MPC, we refer the reader to [10].

Typically MPC is carried out over a finite field and therefore it is only shown how to securely carry out summation and multiplication. We would like to compute over the real numbers and therefore we need to be able to carry out other computations as well. Some works embed secure computations on real numbers into finite field operations by representing the real numbers as fixed or floating point numbers. This is for instance the case in the papers [1, 6, 7, 11]. One of the main challenges this approach faces is to carry out real number divisions in a secure way as division in a finite field and integer division are completely different. As an example the division $5/2$ is 8 in the field \mathbb{F}_{11} but it is 2 as integer division. Hence one needs to introduce other tools to carry out integer division in the finite field. For instance, [1, 7] relies on an iterative algorithm for computing the division. But this comes with the drawback that they need 17 communication rounds before they obtain the result if they use bitlength 32 for their floating points or fixed points. Instead of transforming the real number computations into finite field operations we take another approach and consider the computations directly in the real numbers.

To obtain simpler protocols for the secure arithmetic over the real numbers, we consider real number secret sharing introduced in [22]. This concept is inspired by Shamir’s secret sharing over finite fields [21], and considering this in the real numbers comes with the drawback that a share might reveal a small amount of information. To secretly share a value s in a finite field $\in \mathbb{F}$ using Shamir’s scheme, we choose $c_i \in \mathbb{F}$, uniformly at random for $i = 1, 2, \dots, t$ and compute the polynomial

$$f_s(x) = s + c_1x + c_2x^2 + \dots + c_t x^t. \quad (3)$$

The shares are evaluation points on this polynomial, meaning that for an $\alpha \in \mathbb{F}$ a share is $f_s(\alpha)$. The privacy relies on the uniform distribution in \mathbb{F} combined with Lagrange interpolation arguments stating that $t + 1$ points uniquely determines a degree- t polynomial. Hence, we also have the reconstruction property.

C. MPC based on real number secret sharing

We give a brief introduction to the ideas from [22]. One of the obstacles they had was that if the shares were constructed from the polynomial in (3) the information leakage from the shares depends on how close the evaluation point is to 0. Furthermore, there is no uniform distribution on the real axis.

To circumvent the first obstacle the polynomial f_s is constructed in another way by using the Lagrange polynomial as a basis for all polynomial of degree at most t instead of the basis $\{1, x, \dots, x^t\}$. To be a bit more formal, let $\mathcal{P} = \{\alpha_i\}_{i=1, \dots, n}$. Each time a secret s needs to be shared a polynomial is constructed in the following way. Choose randomly a subset $\mathcal{P}'_s \subset \mathcal{P}$ of size t and let $\mathcal{P}_s = \mathcal{P}'_s \cup \{0\}$. Consider the $t + 1$ points $\mathcal{T} = \{(\alpha_i, \beta_i)\}_{i \in \mathcal{P}_s}$, where we use $(\alpha_0, \beta_0) = (0, s)$. The β_i 's for $i > 0$ are chosen randomly according to a normal distribution with variance σ_β^2 . Define the Lagrange basis polynomials as $L_i(x) = \prod_{j \in \mathcal{P}_s \setminus \{i\}} \frac{x - \alpha_j}{\alpha_i - \alpha_j}$. Then we define

$$f_s(x) = \sum_{i \in \mathcal{P}_s} \beta_i L_i(x). \quad (4)$$

Since $L_i(\alpha_j) = 0$ when $i \in \mathcal{T} \setminus \{i\}$ and $L_i(\alpha_i) = 1$, we obtain that $f_s(x)$ is a degree- t polynomial passing through the points in \mathcal{T} . The shares are evaluation points for this polynomial in the elements in \mathcal{P} meaning that a share of the secret s to the i 'th party is $f_s(\alpha_i)$. To denote that s is secretly shared we write $[s]$, which represents the vector $(f_s(\alpha_1), f_s(\alpha_2), \dots, f_s(\alpha_n))$ where the i 'th party knows the i 'th entry.

We end this section with a description of how to compute on the shares. In [22], protocols for securely computing sums, multiplications, and divisions over the real numbers are presented which are generalizations of protocols for finite fields described in [10]. However, we take fairly different approaches than in [22] for multiplications and inversions. The reason for changing the secure multiplication is due to the fact that we can lower the amount of information sent if we assume that $t < \frac{n}{2}$ and take advantage of some of the properties the Shamir secret sharing possess. The reason we change the secure inversion is to correct a small flaw in the privacy analysis in [22]. Furthermore, we mention that privacy analysis for a single operation (one multiplication, one inversion, etc.) is provided in [22], and we will do the same for our secure operations, but how to analyse the privacy of an algorithm as the one we will present in Section III is not described. In this paper, we suggest to use a conditional independence test for that.

To carry out a secure sum, the parties locally add up the shares for the two secrets. Specifically, assume that

$$\begin{aligned} [x] &= (f_x(\alpha_1), f_x(\alpha_2), \dots, f_x(\alpha_n)) \\ [y] &= (f_y(\alpha_1), f_y(\alpha_2), \dots, f_y(\alpha_n)) \end{aligned} \quad (5)$$

then by $[x] + [y]$ we mean that the parties locally sum up their shares. We remark that $f_x(0) = x$ and $f_y(0) = y$ and the sum of the shares represents a degree- t polynomial f_{x+y} which evaluates to $x + y$ at zero. Hence, we can write $[x] + [y] =$

$[x + y]$. Similarly, we can add a known value to a secret shared value by letting the parties add the value to their share and we can multiply a known value to a secret shared value by letting each party multiply its share by that given value. Thus, we also use the notation $a + [x] = [a + x]$ and $a \cdot [x] = [a \cdot x]$ to represent these computations.

We remark that all operations so far have not required any communication between the parties if $[x]$ and $[y]$ are obtained beforehand. However, for the multiplication and division protocols we require some communication.

For the secure multiplication we assume that $t < \frac{n}{2}$ which allows us to carry out a single multiplication (and linear combinations of such values) locally. Notice that $(f_x(\alpha_1)f_y(\alpha_1), f_x(\alpha_2)f_y(\alpha_2), \dots, f_x(\alpha_n)f_y(\alpha_n))$ corresponds to evaluations of a degree- $2t$ polynomial with constant term xy . Thus a multiplication can also be carried out locally and we write $[x] \cdot [y] = [xy]_{2t}$. Linear combination of $[\cdot]$ and $[\cdot]_{2t}$ shares can be carried out locally as well but if $[\cdot]_{2t}$ shares needs to be multiplied by another secret value we need to "refresh" the shares. To refresh $[x]_{2t}$ we need some preprocessed shares $([r]_{2t}, [r])$, where r is normal distributed with variance σ_r^2 . Then, $[x]_{2t} - [r]_{2t}$ can be computed and opened, and afterwards $x - r + [r] = [x]$. When we say a value is opened it means that all parties reveal their share of the value and the value can be reconstructed. In practice, and in our analysis later on, we implement this as all parties are sending their share to the first party who reconstruct the value and send it back.

For the multiplication, we converted $[x]_{2t}$ to a degree- t sharing. In the remaining we skip the subscript $2t$ for simplicity but we note that one has to take care of the conversions in the algorithm. We show that revealing $x - r$ do not give much information about x if σ_r^2 is high.

2 Proposition:

Let x and r be independent and let $r \sim \mathcal{N}(\mu_r, \sigma_r^2)$ and $\sigma_x^2 = \text{Var}(x)$. Then

$$I(x; x - r) \leq \frac{1}{2} \log \left(1 + \frac{\sigma_x^2}{\sigma_r^2} \right).$$

Proof: We use the definition of mutual information

$$I(x; x - r) = h(x - r) - h(x - r|x).$$

The last term is the entropy of a normal distribution and the first term can be upper bounded by the entropy of such a distribution since for a given variance the entropy of a normal distribution is maximal. Hence, we have

$$\begin{aligned} I(x; x - r) &\leq \frac{1}{2} \log(2\pi e(\sigma_x^2 + \sigma_r^2)) - \frac{1}{2} \log(2\pi e\sigma_r^2) \\ &= \frac{1}{2} \log \left(1 + \frac{\sigma_x^2}{\sigma_r^2} \right). \end{aligned}$$

At last, we describe how to invert real numbers securely. Here, we need some preprocessed $[be^r]$, where r is normal distributed with variance σ_e^2 and b is chosen

uniform randomly from $\{-1, 1\}$. First $[x] \cdot [be^r] = [xbe^r]_{2t}$ can be computed locally as described above and opened towards the parties. After xbe^r is opened we can compute $(xbe^r)^{-1}[be^r] = [x^{-1}]$. We remark that this secure inversion is a bit different than in [22] since we assume another distribution on the obfuscation. Therefore, we prove that a single inversion is not leaking much information.

I HAVE CHANGED THIS PART A BIT. WE MIGHT NEED TO DO SOMETHING IN EXPERIMENTS.

3 Proposition:

Let x and r be independent and let $r \sim \mathcal{N}(\mu_{e^r}, \sigma_{e^r}^2)$ and $\sigma_z^2 = \text{Var}(\log|x|)$. Then

$$I(x; xbe^r) \leq \frac{1}{2} \log \left(1 + \frac{\sigma_z^2}{\sigma_{e^r}^2} \right).$$

Proof: We use that $I(x; y, z) = I(x; g(y, z))$ when g bijective [14], meaning that

$$I(x; xbe^r) = I(x; \text{sgn}(xb), \log|x| + r).$$

Note that $\text{sgn}(xb)$ is -1 with probability $\frac{1}{2}$ and 1 with probability $\frac{1}{2}$ no matter the distribution on x . Hence, $\text{sgn}(xb)$ is independent on x and $\log|x| + r$ and we can remove this term. Using the definition of mutual information we obtain

$$I(x; \log|x| + r) = h(\log|x| + r) - h(\log|x| + r|x).$$

The last term is the entropy of a normal distribution and the first can be upper bounded as in the proof of Proposition 2.

$$\begin{aligned} I(x; xbe^r) &\leq \frac{1}{2} \log(2\pi e(\sigma_z^2 + \sigma_{e^r}^2)) - \frac{1}{2} \log(2\pi e\sigma_{e^r}^2) \\ &= \frac{1}{2} \log \left(1 + \frac{\sigma_z^2}{\sigma_{e^r}^2} \right). \end{aligned}$$

■

Remark that in both Propostion 2 and 3 we can control the leakage of x by varying the variance of r .

We will slightly abuse the notation and write $[X]$ and $[\mathbf{x}]$ for a matrix X and vector \mathbf{x} . This means that all entries should be secret shared and the computations are entry-wise.

D. PAC Bayesian and linear regression

Section II-A considered the loss of a given model. In this section we want to determine the model from a Bayesian point of view, and hence we assume that f is random.

We consider the generalization error from [2] and [13] which states that with probability more than $1 - \delta$ the following inequality is satisfied

$$\begin{aligned} E_{f \sim \hat{\rho}} \mathcal{L}(f) &\leq E_{f \sim \hat{\rho}} \mathcal{L}_N(f) \\ &+ \frac{1}{\lambda} \left(KL(\hat{\rho}, \pi) + \ln \frac{1}{\delta} + \Psi_{\ell, \pi}(\lambda, N) \right), \end{aligned} \quad (6)$$

where π is the prior distribution of f , $\hat{\rho}$ the posterior distribution, λ is a tuning parameter which we return to, and $\Psi_{\ell, \pi}(\lambda, N) = \ln E_{f \sim \pi} E_{(\mathbf{x}, y)} e^{\lambda(\mathcal{L}(f) - \mathcal{L}_N(f))}$, where $\mathcal{L}(f)$ and $\mathcal{L}_N(f)$ are as in (1) and below this equation.

The left-hand side in (6) is the expected generalization error, so we upper bound this with a high probability. For

a given data set, tuning parameters and prior, in order to minimize the right-hand side we need a posterior $\hat{\rho}$ minimizing $E_{f \sim \hat{\rho}} \mathcal{L}_N(f) + \frac{1}{\lambda} KL(\hat{\rho}, \pi)$ since these are the only terms depending on $\hat{\rho}$. The Gibbs posterior

$$\hat{\rho}^*(f) = \frac{1}{E_{f \sim \pi} e^{-\lambda \mathcal{L}_N(f)}} \pi(f) e^{-\lambda \mathcal{L}_N(f)} \quad (7)$$

minimizes this [2]. To choose f from $\hat{\rho}^*(f)$, we set $f_* = \text{argmax}_f \hat{\rho}^*(f)$. Hence, w.r.t. f , we need to minimize

$$\lambda \mathcal{L}_N(f) - \ln \pi(f). \quad (8)$$

We define our model space \mathcal{F} to consists of linear models $f(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x}$ for some $\mathbf{w} \in \mathbb{R}^d$. We put some restrictions on our prior. Since f is uniquely determined by \mathbf{w} we define the prior on \mathbf{w} . We assume a normal distributed prior $\pi(\mathbf{w}) = \mathcal{N}(\boldsymbol{\mu}_w; \Sigma_w)$, where Σ_w is positive definite. With these assumptions (8) reduces to that we have to minimize

$$\frac{\lambda}{N} \sum_{i=1}^n \sum_{j=1}^{n_i} (\mathbf{w} \cdot \mathbf{x}_{ij} - y_{ij})^2 - \frac{1}{2} (\mathbf{w} - \boldsymbol{\mu}_w)^T \Sigma_w^{-1} (\mathbf{w} - \boldsymbol{\mu}_w)$$

with respect to \mathbf{w} . This is a convex function, so differentiating and setting equal to $\mathbf{0}$ yields the minimum. Hence, to minimize it, we solve the following equation.

$$\left(\frac{2\lambda}{N} \sum_{i=1}^n \sum_{j=1}^{n_i} \mathbf{x}_{ij} \mathbf{x}_{ij}^T + \Sigma_w^{-1} \right) \mathbf{w} = \frac{2\lambda}{N} \sum_{i=1}^n \sum_{j=1}^{n_i} \mathbf{x}_{ij} y_{ij} + \Sigma_w^{-1} \boldsymbol{\mu}_w. \quad (9)$$

Setting $\lambda = 0$ we obtain $\mathbf{w} = \boldsymbol{\mu}_w$ meaning that the posterior relies fully on the prior. Opposite, letting λ tend to infinity the posterior relies fully on the data.

III. SECURE DISTRIBUTED REGRESSION USING PAC BAYES APPROACH

Consider the distributed part, where the parties want to find the solution to (9) without revealing the datapoints in \mathcal{D}_i . Note that the parties can compute the inner sums locally,

$$X_i = \sum_{j=1}^{n_i} \mathbf{x}_{ij} \mathbf{x}_{ij}^T, \quad \mathbf{z}_i = \sum_{j=1}^{n_i} \mathbf{x}_{ij} y_{ij},$$

implying that (9) can be rewritten as

$$\left(\frac{2\lambda}{N} \sum_{i=1}^n X_i + \Sigma_w^{-1} \right) \mathbf{w} = \frac{2\lambda}{N} \sum_{i=1}^n \mathbf{z}_i + \Sigma_w^{-1} \boldsymbol{\mu}_w.$$

This is similar to the local part in [19], but we remark that the setup in their paper is different from ours since the computation takes place at some servers.

We need now a secure way to compute the matrix $A = \frac{2\lambda}{N} \sum_{i=1}^n X_i + \Sigma_w^{-1}$ and the vector $\mathbf{b} = \frac{2\lambda}{N} \sum_{i=1}^n \mathbf{z}_i + \Sigma_w^{-1} \boldsymbol{\mu}_w$. Furthermore, we need a way to solve $A\mathbf{w} = \mathbf{b}$ in a secure way without revealing A and \mathbf{b} . We start with the i 'th party locally compute X_i and \mathbf{z}_i and secretly share the entries between all parties. Then by using the computations on shares described in Section II-C the parties can compute $[\frac{2\lambda}{N} \sum_{i=1}^n X_i + \Sigma_w^{-1}]$ and $[\frac{2\lambda}{N} \sum_{i=1}^n \mathbf{z}_i + \Sigma_w^{-1} \boldsymbol{\mu}_w]$. Now we have A and \mathbf{b} secretly shared and we need a way to secretly solve the system $A\mathbf{w} = \mathbf{b}$.

Input: P_i holds $\mathcal{D}_i = \{(\mathbf{x}_{ij}, y_{ij}) \mid j = 1, 2, \dots, n_i\}$, where $N = \sum_{i=1}^n n_i$ is known. The tuning parameter λ , $\boldsymbol{\mu}_w$, and a positive definite Σ_w determining the prior on \mathbf{w} is fixed.

Output: The parties should learn \mathbf{w} satisfying (9).

The protocol:

- 1) P_i sets $X_i = \sum_{j=1}^{n_i} \mathbf{x}_{ij} \mathbf{x}_{ij}^T$ and $\mathbf{z}_i = \sum_{j=1}^{n_i} \mathbf{x}_{ij} y_{ij}$
- 2) P_i send shares of X_i and \mathbf{z}_i to the other parties.
- 3) Using the secure protocols for summation, multiplication by scalar, and adding scalar, the parties compute

$$[A] = \frac{2\lambda}{N} \sum_{i=1}^n [X_i] + \Sigma_w^{-1}$$

$$[\mathbf{b}] = \frac{2\lambda}{N} \sum_{i=1}^n [\mathbf{z}_i] + \Sigma_w^{-1} \boldsymbol{\mu}_w$$

- 4) Denote by $[C] = [[A] \mid [\mathbf{b}]$ the $d \times (d+1)$ total matrix with shared entries $[c_{ij}]$
- 5) **for** $k = 1 \dots d - 1$ **do**
 $[\text{inv}_k] = \frac{1}{[c_{kk}]}$
for $i = k + 1 \dots d$ **do**
 $[\text{frac}_{ik}] = [c_{ik}] \cdot [\text{inv}_k]$
 $[c_{ik}] = [0]$
for $j = k + 1 \dots d + 1$ **do**
 $[c_{ij}] = [c_{ij}] - [\text{frac}_{ik}] \cdot [c_{kj}]$

C is now an upper triangular matrix with nonzero elements on the diagonal

- 6) **for** $i = 0 \dots d - 1$ **do**
 $[w_{d-i}] = \frac{[c_{d-i, d+1}] - \sum_{j=1}^i [c_{d-i, d-i+j}] w_{d-i+j}}{[c_{d-i, d-i}]}$
Open w_{d-i} to all parties

Protocol 1. Secure Regression with prior from Gaussian elimination

Notice that X_i is a positive semidefinite matrix and hence A is positive definite due to the assumption on Σ_w . Thus, the matrix is invertible and the system in (9) has a unique solution. Furthermore, since A is symmetric and positive definite we can do Gaussian elimination without pivoting. This means that we never have to interchange rows when row-reducing since the diagonal entries will always be nonzero. After we obtain an upper-triangular matrix, we solve the system with backward substitution. The protocol can be found in Protocol 1.

A. Communication Analysis

We analyse the amount of information the parties need to send between each other depending on the size of the matrix in Protocol 1. We remark that the only it is only shares, multiplications, inversions and opening of shared values which require that parties need to send information. First party i has to construct shares for X_i and \mathbf{z}_i . This require him to send $(d^2 + d)$ values to the $(n-1)$ other parties. Furthermore note that both multiplication and inversion also consists of openings. In worst case both a multiplication and

an inversion costs 2 openings.⁴

For the Gaussian elimination part, we use $d-1$ inversions for computing $[\text{inv}_k]$ and $\frac{d^2-d}{2}$ multiplications to compute $[\text{frac}_{ik}]$. Then we use $\frac{d^3-d}{3}$ multiplications in the inner for-loop, and at last we do d inversions, multiplications, and openings in the last for-loop. This gives a total of $\frac{2d^3+3d^2+d}{6}$ multiplications, $2d-1$ inversions, and additional d openings. In conclusion, we get that the total amount of openings are at most $\frac{2d^3+3d^2+16d-6}{3}$.

IV. EXPERIMENTS

A. Data set and experimental setup

In our experiments, we used Boston house price data set, a copy of UCI ML housing data set, that is contained in sklearn library [12, 20], and which is suitable for linear regression. This data set contains 506 instances with 13 feature targeting the median value of the prices of the homes in \$1000's. Before we applied our models to this data, we normalized it between 0 and 1.

B. Testing accuracy

In our experiments, we used the mean square error (MSE) to evaluate the accuracy of our secure regression method. To make our experiments trustworthy, we held 10 differently evaluations for each parameter combination. The data set is divided into train set (80%) and test set (20%). We shuffle the data set for each 10 evaluations. We divide the train set into n subsets, one for each party, whenever we use secure computations for privacy preserving algorithms as mentioned in Section II-A. Mean MSE results are reported across 10 different evaluations along with the corresponding standard deviations for the Gaussian elimination method with and without secure computation.

In Table I, we report the mean square error (MSE) for each algorithm using different λ values from the set $\{0.01, 0.1, 1, 10, 100, 1000\}$. We see that the secure version are very similar to the insecure one with respect to accuracy. However, we experience that increase in t and n may produce deviations in the MSE score for secure Gaussian elimination. For example, for SGE with $n = 9, t = 4$, we observe deviations for large λ 's.

In Table II, we investigate the effect of the security parameters, σ_r^2 and σ_β^2 , on the accuracy-privacy trade-off. Due to numerical problems, e^r gave some NaN when r was negative. Furthermore, very large r was problematic here as well. Thus we have fixed $\mu_{e^r} = 100$ and $\sigma_{e^r}^2 = 400$. We observe that as we increase σ_r^2 and σ_β^2 , our secure computations become unstable and result in larger deviations from the right values.

⁴multiplication if both values needs to be converted to a degree- t sharing and an inversion since it consists of a product where only the value x we want to invert might need an conversion to degree- t and then an opening of xbe^r

λ	0.01	0.1	1	10	100	1000
GE	0.17014 \pm 0.02350	0.08442 \pm 0.01782	0.03167 \pm 0.00818	0.01737 \pm 0.00384	0.01556 \pm 0.00238	0.01614 \pm 0.00224
SGE ($n = 3, t = 1$)	0.17014 \pm 0.02350	0.08442 \pm 0.01782	0.03167 \pm 0.00818	0.01737 \pm 0.00384	0.01556 \pm 0.00238	0.01614 \pm 0.00224
SGE ($n = 5, t = 2$)	0.17014 \pm 0.02727	0.08442 \pm 0.01782	0.03167 \pm 0.00818	0.01737 \pm 0.00384	0.01556 \pm 0.00238	0.01614 \pm 0.00224
SGE ($n = 7, t = 3$)	0.17014 \pm 0.02350	0.08442 \pm 0.01782	0.03167 \pm 0.00817	0.01737 \pm 0.00384	0.01556 \pm 0.00238	0.01615 \pm 0.00224
SGE ($n = 9, t = 4$)	0.17016 \pm 0.02351	0.08442 \pm 0.01782	0.03168 \pm 0.00817	0.03370 \pm 0.03224	0.04106 \pm 0.07312	0.02682 \pm 0.03297

TABLE I

MSE RESULTS FOR DIFFERENT λ WITH PRIOR $\Sigma_w = I$ AND $\mu_w = \mathbf{0}$. WE USE SECURITY PARAMETERS $\sigma_r^2 = 10^8, \sigma_\beta^2 = 10^8$

λ	0.01	0.1	1	10	100	1000
SGE ($\sigma_r^2 = 10^8, \sigma_\beta^2 = 10^8$)	0.17014 \pm 0.02727	0.08442 \pm 0.01782	0.03167 \pm 0.00818	0.01737 \pm 0.00384	0.01556 \pm 0.00238	0.01614 \pm 0.00224
SGE ($\sigma_r^2 = 10^{10}, \sigma_\beta^2 = 10^{10}$)	0.17014 \pm 0.02350	0.08442 \pm 0.01781	0.03167 \pm 0.00818	0.01737 \pm 0.00384	0.01555 \pm 0.00239	0.01615 \pm 0.00223
SGE ($\sigma_r^2 = 10^{12}, \sigma_\beta^2 = 10^{12}$)	0.17014 \pm 0.02350	0.08442 \pm 0.01782	0.03166 \pm 0.00816	0.01733 \pm 0.00374	0.01543 \pm 0.00249	0.01614 \pm 0.00224
SGE ($\sigma_r^2 = 10^{14}, \sigma_\beta^2 = 10^{14}$)	0.17014 \pm 0.02350	0.08442 \pm 0.01783	0.03171 \pm 0.00823	0.01751 \pm 0.00390	0.39457 \pm 0.86603	0.45492 \pm 1.18378

TABLE II

MSE RESULTS FOR DIFFERENT SECURITY PARAMETERS WHEN $n = 5$ AND $t = 2$. FIRST ROW IS THE SAME ROW AS ($n = 5, t = 2$) IN TABLE I

sample size (m):	200	300	400	500
$\sigma_\beta^2 = 10^{-8}, \sigma_r^2 = 10^{-6}, \sigma_e^2 = 10^{-6}, \mu_{e^r} = 1$	0.194 \pm 0.132	0.040 \pm 0.089	0.021 \pm 0.034	0.013 \pm 0.038
$\sigma_\beta^2 = 100, \sigma_r^2 = 100, \sigma_e^2 = 16, \mu_{e^r} = 10$	0.233 \pm 0.184	0.131 \pm 0.147	0.014 \pm 0.022	0.015 \pm 0.031
$\sigma_\beta^2 = 10^8, \sigma_r^2 = 10^8, \sigma_e^2 = 400, \mu_{e^r} = 100$	0.147 \pm 0.090	0.193 \pm 0.167	0.103 \pm 0.166	0.024 \pm 0.018
$\sigma_\beta^2 = 10^{12}, \sigma_r^2 = 10^{12}, \sigma_e^2 = 400, \mu_{e^r} = 100$	0.262 \pm 0.230	0.246 \pm 0.175	0.091 \pm 0.153	0.026 \pm 0.036

TABLE III

MEAN PLUS/MINUS STANDARD DEVIATION OF p -VALUES FOR THE 10 CONDITIONAL INDEPENDENCE TESTS WITH m SAMPLES.

C. Privacy analysis – Boston data set

We consider the privacy of our presented protocol in a 3-party setting, where the computation is on the Boston data set. To analyze the privacy, we make use of the conditional independence test from [3] to argue that the conditional mutual information in (2) is small. We assume that the third party is corrupt and hence V consists of; the randomness he chooses when secret sharing the entries in X_3 and \mathbf{z}_3 , his shares of the preprocessed values for multiplications and divisions, the shares he receives for X_i and $\mathbf{z}_i, i = 1, 2$, and the opened values during the protocol. We denote this V by V_3 since it is the view of the third party and hence with this setup (2) becomes

$$I(X_1, X_2, \mathbf{z}_1, \mathbf{z}_2; V_3 | X_3, \mathbf{z}_3, \mathbf{w}).$$

For simplicity and for illustration of how the security parameters impact the privacy we fixed $\alpha_i = i$ and $\mathcal{P}'_s = \{\alpha_1\}$, when generating the shares. We remark that intuitively it will be more secure to choose the elements in \mathcal{P}'_s randomly but this seems to confuse the test to not reject the null hypothesis even with small security parameters when we do not have that many samples. So for illustration, we decided to fix \mathcal{P}'_s .

We want to test the null hypothesis that $X_1, X_2, \mathbf{z}_1, \mathbf{z}_2$ and V_3 are conditionally independent given $X_3, \mathbf{z}_3, \mathbf{w}$. Of course this is not the case since V_3 for instance includes shares for $X_1, X_2, \mathbf{z}_1, \mathbf{z}_2$, but with high security parameters, the dependency should be masked implying that it is close to be conditionally independent. Hence $I(X_1, X_2, \mathbf{z}_1, \mathbf{z}_2; V_3 | X_3, \mathbf{z}_3, \mathbf{w})$ must be close to zero. To illustrate this, we apply the test from [3] and we need samples of $X_1, X_2, \mathbf{z}_1, \mathbf{z}_2, V_3, X_3, \mathbf{z}_3, \mathbf{w}$. We construct such

samples in the following way. For a single sample, we choose randomly 20% of the data instances from the Boston data set and split this subset in three, one for each party. From this subset we obtain a sample for $X_1, X_2, X_3, \mathbf{z}_1, \mathbf{z}_2, \mathbf{z}_3$. These are all the inputs to our secure protocol (Protocol 1), and running this protocol will give a sample for V_3 and \mathbf{w} . For some fixed security parameters ($\sigma_\beta^2, \sigma_r^2, \sigma_{e^r}^2$) this process is repeated 2000 times to obtain 2000 samples of the random vectors. We feed the conditional independence test 10 times with m randomly chosen of these samples. We report the mean and standard deviation of the p -value in Table III. As we observe, the conditional independence test needs more samples before it can determine any conditional dependencies, and reject the null hypothesis, when we are increasing the security parameters.

V. CONCLUSION AND FUTURE WORK

We have described a secure protocol for obtaining a linear model based on a prior and on a number of data sets distributed between several parties. We have illustrated that carrying out the computations using MPC techniques based on real number secret sharing we can obtain the model and argued that this technique is secure even when multiple secure operations are carried out sequentially. To argue security we showed that by increasing the security parameters it is hard to determine any dependency between the information we want to hide and the messages received during the protocol. However, if the security parameters become too high, we observed some inaccuracy in the secure Gaussian elimination protocol due to numerical problems. In this paper, we have focused on how real number MPC can be used to obtain a linear model on a split data set but it could be interested to consider other models as well. However, if the model can be obtained from linear combinations, products and inversions we have in this paper the tools to obtain the model even though further analysis of the privacy is of course needed. Furthermore, it could be interesting to consider other secure operations to expand the toolbox for secure computations.

REFERENCES

- [1] Mehrdad Aliasgari, Marina Blanton, Yihua Zhang and Aaron Steele. “Secure Computation on Floating Point Numbers”. In: *NDSS*. 2013.
- [2] Pierre Alquier, James Ridgway and Nicolas Chopin. “On the properties of variational approximations of Gibbs posteriors”. In: *Journal of Machine Learning Research* 17.236 (2016), pp. 1–41. URL: <http://jmlr.org/papers/v17/15-290.html>.
- [3] Alexis Bellot and Mihaela van der Schaar. “Conditional Independence Testing using Generative Adversarial Networks”. In: *NeurIPS*. 2019.
- [4] Raphael Bost, Raluca Ada Popa, Stephen Tu and Shafi Goldwasser. “Machine learning classification over encrypted data”. In: *Cryptology ePrint Archive* (2014).
- [5] Olivier Catoni. “Pac-Bayesian Supervised Classification: The Thermodynamics of Statistical Learning”. In: *IMS Lecture Notes Monograph Series* (Jan. 2008). DOI: 10.1214/074921707000000391.
- [6] Octavian Catrina and Claudiu Dragulin. “Multiparty Computation of Fixed-Point Multiplication and Reciprocal”. In: *20th International Workshop on Database and Expert Systems Application*. 2009. DOI: 10.1109/DEXA.2009.84.
- [7] Octavian Catrina and Amitabh Saxena. “Secure Computation with Fixed-Point Numbers”. In: *Financial Cryptography and Data Security*. Jan. 2010, pp. 35–50. DOI: 10.1007/978-3-642-14577-3_6.
- [8] Martine de Cock, Rafael Dowsley, Anderson C.A. Nascimento and Stacey C. Newman. “Fast, Privacy Preserving Linear Regression over Distributed Datasets Based on Pre-Distributed Data”. In: *Proceedings of the 8th ACM Workshop on Artificial Intelligence and Security*. 2015, 3–14. DOI: 10.1145/2808769.2808774.
- [9] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. Wiley-Interscience, 2006. ISBN: 0471241954.
- [10] Ronald Cramer, Ivan Bjerre Damgård and Jesper Buus Nielsen. *Secure Multiparty Computation and Secret Sharing*. Cambridge University Press, 2015. DOI: 10.1017/CBO9781107337756.
- [11] Vassil Dimitrov, Liisi Kerik, Toomas Krips, Jaak Randmets and Jan Willemsen. “Alternative Implementations of Secure Real Numbers”. In: *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*. Association for Computing Machinery, 2016. DOI: 10.1145/2976749.2978348.
- [12] Dheeru Dua and Casey Graff. *UCI Machine Learning Repository*. 2017. URL: <https://archive.ics.uci.edu/ml>.
- [13] Pascal Germain, Francis Bach, Alexandre Lacoste and Simon Lacoste-Julien. “PAC-Bayesian Theory Meets Bayesian Inference”. In: *Proceedings of the Neural Information Processing Systems Conference*. 2016. URL: <https://hal.science/hal-01324072>.
- [14] Alexander Kraskov, Harald Stögbauer and Peter Grassberger. “Estimating mutual information”. In: *Phys. Rev. E* 69 (6 2004). DOI: 10.1103/PhysRevE.69.066138.
- [15] Bulut Kuskonmaz, Jaron S. Gundersen and Rafal Wisniewski. “Investigation of Alternative Measures for Mutual Information”. In: *IFAC-PapersOnLine* 55.16 (2022). 18th IFAC Workshop on Control Applications of Optimization CAO 2022, pp. 154–159. DOI: 10.1016/j.ifacol.2022.09.016.
- [16] John Langford and Matthias Seeger. *Bounds for Averaging Classifiers*. 2001.
- [17] Qiongxiu Li, Jaron Skovsted Gundersen, Richard Heusdens and Mads Græsbøll Christensen. “Privacy-Preserving Distributed Processing: Metrics, Bounds and Algorithms”. In: *IEEE Transactions on Information Forensics and Security* 16 (2021), pp. 2090–2103. DOI: 10.1109/TIFS.2021.3050064.
- [18] David McAllester. “Simplified PAC-Bayesian Margin Bounds”. In: *Learning Theory and Kernel Machines*. Springer Berlin Heidelberg, 2003, pp. 203–215.
- [19] Valeria Nikolaenko, Udi Weinsberg, Stratis Ioannidis, Marc Joye, Dan Boneh and Nina Taft. “Privacy-Preserving Ridge Regression on Hundreds of Millions of Records”. In: *IEEE Symposium on Security and Privacy*. 2013, pp. 334–348. DOI: 10.1109/SP.2013.30.
- [20] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot and E. Duchesnay. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.
- [21] Adi Shamir. “How to Share a Secret”. In: *Commun. ACM* 22.11 (1979), 612–613. DOI: 10.1145/359168.359176.
- [22] Katrine Tjell and Rafael Wisniewski. *Privacy in Distributed Computations based on Real Number Secret Sharing*. 2021. arXiv: 2107.00911 [cs.CR].
- [23] Kun Zhang, Jonas Peters, Dominik Janzing and Bernhard Schölkopf. “Kernel-Based Conditional Independence Test and Application in Causal Discovery”. In: *AUAI Press*, 2011, 804–813. ISBN: 9780974903972.