

Adrià Gascón, Phillipp Schoppmann, Borja Balle, Mariana Raykova, Jack Doerner, Samee Zahur, and David Evans

Privacy-Preserving Distributed Linear Regression on High-Dimensional Data

Abstract: We propose privacy-preserving protocols for computing linear regression models, in the setting where the training dataset is vertically distributed among several parties. Our main contribution is a hybrid multi-party computation protocol that combines Yao’s garbled circuits with tailored protocols for computing inner products. Like many machine learning tasks, building a linear regression model involves solving a system of linear equations. We conduct a comprehensive evaluation and comparison of different techniques for securely performing this task, including a new Conjugate Gradient Descent (CGD) algorithm. This algorithm is suitable for secure computation because it uses an efficient fixed-point representation of real numbers while maintaining accuracy and convergence rates comparable to what can be obtained with a classical solution using floating point numbers. Our technique improves on Nikolaenko et al.’s method for privacy-preserving ridge regression (S&P 2013), and can be used as a building block in other analyses. We implement a complete system and demonstrate that our approach is highly scalable, solving data analysis problems with one million records and one hundred features in less than one hour of total running time.

DOI 10.1515/popets-2017-0053

Received 2017-02-28; revised 2017-06-01; accepted 2017-06-02.

Adrià Gascón: The Alan Turing Institute and University of Warwick, agascon@turing.ac.uk. Part of this work was supported by The Alan Turing Institute under the EPSRC grant EP/N510129/1. Part of this work was done while the author was at the University of Edinburgh, supported by the SOCIAM Project under EPSRC grant EP/J017728/2.

Phillipp Schoppmann: Humboldt University of Berlin, schoppmann@informatik.hu-berlin.de. Supported by the German Research Foundation (DFG) through Research Training Group GRK 1651 (SOAMED). Part of this work was done while the author was at the University of Edinburgh.

Borja Balle: Amazon, pigem@amazon.co.uk. Work done while the author was at Lancaster University.

Mariana Raykova: Yale University, mariana.raykova@yale.edu. Supported by NSF grants CNS-1633282, 1562888, 1565208, and DARPA SafeWare W911NF-15-C-0236, W911NF-16-1-0389.

Jack Doerner: Northeastern University, jhd3pa@virginia.edu. Work done while the author was at the University of Virginia.

Samee Zahur: Google, szahur@google.com

David Evans: University of Virginia, evans@virginia.edu. Supported in part by the National Science Foundation (NSF Award 1111781), Google, and Amazon.

1 Introduction

Predictive modeling is an essential tool in decision making processes in domains such as policy making, medicine, law enforcement, and finance. To obtain more accurate models, it is common that organizations cooperate to build joint training datasets, specially when the analysis at hand involves complex phenomena whose study requires vast amounts of data. However, collaborative analyses involving private data are often limited by ethical and regulatory constraints. As an example, consider a research project consisting of a study to predict medical conditions given data related to socio-economic background. While databases holding medical, judicial, and tax records linkable by common unique identifiers (e.g., social security numbers) exist, they are often held by different institutions.

One solution is to rely on a trusted third-party that is provided with the databases from each party *in the clear*, which then performs the analysis over the data and returns the result to all parties. This approach requires all participants to agree on one party they trust, and poses an obvious privacy threat to the data. Finding effective ways of limiting the amount of trust to be placed on such third-party is an ongoing research challenge being tackled simultaneously from many perspectives.

Our goal is to enable such joint analyses without the need to expose each organization’s sensitive data to any other party. We consider a natural setup where parties have incentives to collaborate in obtaining the parameters of a regression model, but are not willing to disclose any information about their data beyond what is revealed by the learned model. Secure Multi-Party Computation (MPC) protocols provide a general mechanism for multiple parties to jointly compute a function on their private inputs. However, existing out-of-the-box MPC techniques do not yet scale to problem sizes encountered in real-world data analysis applications. The goal of this work is to circumvent such limitations by developing application-specific efficient protocols with better performance than generic MPC techniques.

The motivation to join databases in the medical example above is to build a more accurate model by using a joint training dataset of *higher dimensionality* than any of the original ones. This is known as data analysis on *vertically-partitioned* datasets, as the columns (i.e., features) in the training dataset are distributed among several parties. This is in contrast with

the *horizontally-partitioned* setting, where parties produce a joint dataset by contributing rows (i.e., data points) involving a fixed set of features.

Linear regression is a fundamental machine learning task that fits a linear curve over a set of high-dimensional data points. An important property of this problem is that it can be cast as an optimization problem whose solution admits a closed-form expression. Formally, linear regression can be reduced to solving a system of linear equations of the form $A\theta = b$. Interestingly, solving systems of linear equations is a basic building block of many other machine learning algorithms [54]. Thus, secure MPC solutions for solving linear systems can also be used to develop other privacy-preserving data analysis algorithms.

Contributions. In this paper, we propose a solution for securely computing a *linear regression* model from a vertically-partitioned dataset distributed among an arbitrary number of parties. Our protocols are the result of securely composing a protocol for privately computing the coefficients of the system of equations $A\theta = b$ from distributed datasets (*aggregation phase*), and a protocol for securely solving linear systems (*solving phase*). Our main contributions are as follows:

- Scalable MPC protocols for linear regression on vertically partitioned datasets (Section 3). We design, analyze, and evaluate two hybrid MPC protocols allowing for different trade-offs between running time and the availability of external parties that facilitate the computation.
- A fast solver for high-dimensional linear systems suitable for MPC (Section 5). Our solver relies on a new Conjugate Gradient Descent (CGD) algorithm that scales better, both in terms of running time and accuracy, than the best previous MPC-based alternative.
- We conduct an extensive evaluation of our system on real data from the UCI repository [46] (Section 6.4). The results show that our system can produce solutions with accuracy comparable to those obtained by standard tools without privacy constraints.

Our work is also the first to undertake a formal analysis of the effect different number representations have on the accuracy of privacy-preserving regression protocols. We further propose a thorough data pre-processing pipeline to mitigate accuracy loss (Section 4).

While the basic versions of our protocols are secure against semi-honest adversaries, in Section 7 we discuss extensions of our protocol beyond semi-honest security. We implemented a protocol extension providing stronger security guarantees for the solving phase and observed that it introduces minimal overhead over the semi-honest solution.

Code for our complete system is publicly available under open-source licenses [75].

Related work. Questions of privacy-preserving data mining and private computation of machine learning algorithms have been considered in several works [19, 40, 48, 71], providing theoretical protocols without implementations or practical evaluations of their efficiency. Early implementations of privacy-preserving linear regression protocols [21, 65] either don't use formal threat models, or leak additional information beyond the result of the computation.

Hall et al. [35] were the first to propose a protocol for linear regression with formally defined security guarantees. However, due to the dependence on expensive homomorphic encryption, the resulting system does not scale well to large datasets.

Another implementation of linear regression is given by Bogdanov et al. [6]. They compare multiple methods for solving regression problems, including standard Conjugate Gradient Descent. However, they work in a threat model with three *computing parties* and a *honest majority* among those. Although they allow for more than three *input parties*, this means that they cannot handle the two-party case. Furthermore, their evaluation is limited to problems with at most 10 features.

Nikolaenko et al.'s work [59] on secure computation for linear regression is the one most similar to ours. In particular, their approach also considers a protocol split into aggregation and solving phases implemented using multiple MPC primitives. However, Nikolaenko et al. only consider the horizontally-partitioned setting while we focus on the more challenging case of vertically-partitioned datasets. Our implementation of the solving phase using CGD improves the results obtained using Nikolaenko et al.'s secure Cholesky decomposition in terms of both speed and accuracy for datasets with more than 100 features (experiments in [59] were limited to datasets with at most 20 features). Furthermore, our CGD algorithm can also be used directly to improve other MPC protocols requiring a secure linear system solver [33, 35, 59].

Vertically-partitioned databases have been also studied from the perspective of privacy-preserving querying of statistical databases [22]. This work falls under the paradigm of differential privacy [23], where the goal is to mitigate the privacy risks incurred by revealing to an attacker the model computed by a machine learning algorithm. Our work addresses an orthogonal concern, namely that of preserving privacy during the computation the model. Combining MPC and differential privacy is an interesting problem we plan to address in future work.

2 Background

The following two sections provide background on linear regression and MPC techniques that will be used in this paper.

2.1 Linear Regression

Linear regression is a fundamental building block of many machine learning algorithms. It produces a model given a training set of sample data points by fitting a linear curve through them. Formally, a linear model is a real-valued function on d -dimensional vectors of the form $x \mapsto \langle \theta, x \rangle$, where $x \in \mathbb{R}^d$ is the input vector, $\theta \in \mathbb{R}^d$ is the vector of parameters specifying the model, and $\langle \theta, x \rangle = \sum_{j=1}^d \theta_j x_j$ is the inner product between θ and x . The term *linear* comes from the fact that the output predicted by the function $\langle \theta, x \rangle$ is a linear combination of the features represented in x .

The learning algorithm has access to a *training set* of samples representing the input-output relation the model should try to replicate. These are denoted as $(x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)})$ with $x^{(i)} = (x_1^{(i)}, \dots, x_d^{(i)}) \in \mathbb{R}^d$ and $y^{(i)} \in \mathbb{R}$ for $i \in [n]$.

A standard approach to learn the parameters of a linear model is to solve the *ridge regression* optimization:

$$\operatorname{argmin}_{\theta} \frac{1}{n} \sum_{i=1}^n (y^{(i)} - \langle \theta, x^{(i)} \rangle)^2 + \lambda \|\theta\|^2 \quad (1)$$

Here, $\lambda \|\theta\|^2$ is known as a ridge (or Tikhonov) regularization term weighted by the regularization parameter $\lambda > 0$. When $\lambda = 0$, this optimization finds the parameter vector θ minimizing the *mean squared error* between the predictions made by the model on the training examples $x^{(i)}$ and the desired outputs $y^{(i)}$. The ridge penalty is used to prevent the model from overfitting when the amount of training data is too small.

The optimization problem (1) is convex and its solution admits a closed-form expression. A way to derive this solution is to reformulate the optimization problem by writing $X \in \mathbb{R}^{n \times d}$ for the matrix with n rows corresponding to the different d -dimensional row vectors $x^{(i)\top}$ so that $X(i, j) = x_j^{(i)}$, and $Y \in \mathbb{R}^n$ for a column vector with n entries corresponding to the training labels $y^{(i)}$. With this notation, (1) is equivalent to

$$\operatorname{argmin}_{\theta} \frac{1}{n} \|Y - X\theta\|^2 + \lambda \|\theta\|^2 .$$

Now, by taking the gradient of this objective function and equating it to zero one can see that the optimization (1) reduces to solving this system of linear equations:

$$\left(\frac{1}{n} X^\top X + \lambda I \right) \theta = \frac{1}{n} X^\top Y \quad (2)$$

where I represents the $d \times d$ identity matrix.

Therefore, ridge regression reduces to solving a system of linear equations of the form $A\theta = b$, where the coefficients $A = \frac{1}{n} X^\top X + \lambda I$ and $b = \frac{1}{n} X^\top Y$ only depend on the training data (X, Y) . Since the coefficient matrix A is positive definite by construction ridge regression can be solved using one of the many known algorithms for solving positive definite linear systems.

2.2 Multi-Party Computation

Multi-Party Computation (MPC) is an area of cryptography concerned with enabling multiple parties to jointly evaluate a public function on their private inputs, without leaking any information about their respective inputs (other than their sizes and whatever can be inferred from the result of the computation). Several generic techniques have been developed for MPC including Yao's garbled circuits protocol [70], the GMW protocol [31], and other secret-sharing based techniques [7, 41, 50]. These differ in several aspects including the number of parties they support, the representation they use for the evaluated function (e.g., binary vs. arithmetic circuits), the function families that can be evaluated securely with these solutions (e.g., bounded-degree polynomials vs. arbitrary polynomials), as well trade-offs in terms of communication, computation, and security guarantees that they provide. In Section 3, we keep these differences in mind in order to choose the right MPC techniques for each part of the computation. Popular applications used for the empirical evaluation of practical MPC systems include graph algorithms [5, 32, 42, 51, 57], data structure algorithms [42, 51], string matching and distance algorithms [36, 37, 66] and AES [3, 36, 37, 66].

Other than in Section 7, this paper assumes *semi-honest adversaries*. Such adversaries are limited to observing the full trace of the protocol execution without influencing it. This contrasts with *malicious adversaries*, that take complete control of protocol participants, and may deviate from the protocol in arbitrary ways. For detailed definitions of security in both models see Goldreich [30]. To prove security, we use a simulation-based approach [30, 47] where one argues that a protocol is secure if it can be simulated in an ideal environment in which there is no data leakage by definition.

For our solving phase, we use Yao's two-party garbled circuits protocol, which was formally-described and proven secure by Lindell and Pinkas [49]. In its basic form, this protocol provides semi-honest security for computing arbitrary functions represented as a binary circuit. A two-party computation protocol based on garbled circuits has three phases: circuit garbling, input garbling and garbled circuit evaluation. In the first step one of the parties (the garbler) creates a garbled represen-

tation of the evaluated Boolean circuit where wire values are represented with cryptographic keys (garbled labels) and the circuit gates are implemented as garbled tables of ciphertexts that are encrypted using the wire labels. The garbled circuit consisting of all garbled gates is provided to the other party (the evaluator) who can evaluate it on the input for which it is given the corresponding input garbled labels. These input garbled values are provided in the input garbling phase. While the garbler can directly send the garbled values for its input bits, the parties need to run an oblivious transfer (OT) protocol [55, 62] to enable the evaluator to obtain the input garbled labels corresponding its input bits without revealing those inputs to the generator.

3 Protocol description

We consider the problem of solving a ridge regression problem when the training data ($X \in \mathbb{R}^{n \times d}, Y \in \mathbb{R}^n$) is distributed vertically (by columns) among several parties. As in previous work on privacy-preserving ridge regression [35, 59], we assume that the regularization parameter λ is public to all the parties, and hence the implementation of a secure data analysis pipeline including hyper-parameter tuning (e.g., cross-validation) is beyond the scope of this paper. Furthermore, we also make the assumption in the vertically partitioned case that the correspondence between the rows in the datasets owned by different parties is known a priori. This assumption is easy to enforce in databases with unique identifiers for each record, given that efficient private set intersection protocols exist. The more complex task of privacy-preserving entity resolution is beyond the scope of this paper.

As discussed in Section 2.1, solving a ridge regression problem reduces to the solution of a positive definite linear system. A superficial inspection of Equation (2) shows that this can be seen as two sequential tasks: (i) compute the matrix $A = \frac{1}{n} X^\top X + \lambda I$ and the vector $b = \frac{1}{n} X^\top Y$, and (ii) solve the system $A\theta = b$. In the distributed privacy-preserving setting, task (i) above corresponds to an *aggregation phase* where data held by different parties has to be combined. In contrast, we refer to (ii) as the *solving phase*, where the system of linear equations resulting from the aggregation is solved and all the parties obtain the solution in some form. It is important to note that in practice, ridge regression is used on problems with $n \gg d$, and that, while the input of the aggregation phase is of the order of $n \times d$, the solving phase has input size independent of n , i.e., of the order of d^2 . This demonstrates the importance of having an efficient aggregation phase.

As mentioned in the introduction, the scenario where (X, Y) is partitioned horizontally among the parties was stud-

ied by Nikolaenko et al. [59]. With this partitioning, the computation of A and b can be cast as an addition of matrices computed *locally* by the parties. More concretely, in that case one can write $A = \sum_{i=1}^n A_i + \lambda I$ with $A_i = x^{(i)} x^{(i)\top} / n$, and similarly $b = \sum_{i=1}^n b_i$ with $b_i = x^{(i)} y^{(i)} / n$. The absence of the need for secure multiplication is crucial from a secure computation perspective, and guided the design of the protocol of [59], as non-linear operations are expensive in secure computation. In contrast, the setting where the data is vertically partitioned requires secure computation of inner products on vectors owned by different parties. In the following section, we analyze the exact requirements for this setting, and propose secure computation protocols to overcome this challenge.

3.1 Aggregation Phase

For clarity of presentation, we first describe the two-party case, where X is vertically partitioned among two parties P_1 and P_2 as X_1 and X_2 , and Y is held by one of the parties, say P_2 . As described above, the goal of this phase is to have the parties hold additive shares of (A, b) . With this notation, the functionality f of the aggregation phase can be described as follows:

$$f(X_1, X_2, Y) = ((A - A_2, b - b_2), (A_2, b_2)) ,$$

where $(A - A_2, b - b_2)$ and (A_2, b_2) are the outputs received by the first and the second party respectively.

Our protocol implements f in a secure way. Let us now have a closer look at what needs to be computed. For simplicity, assume that each party holds a block of contiguous features. Then, the following equations show how the content of the output matrix depends on the inputs of the two parties.

$$X^\top = \begin{bmatrix} X_1^\top \\ X_2^\top \end{bmatrix} \quad X = [X_1 \ X_2]$$

$$X^\top X = \begin{bmatrix} X_1^\top X_1 & X_1^\top X_2 \\ X_2^\top X_1 & X_2^\top X_2 \end{bmatrix}$$

Observe that the upper left part of the matrix $M = X^\top X$ depends only on the input of party P_1 and the lower right part depends only on the input of party P_2 . Hence, the corresponding entries of M can be computed locally by P_1 , while P_2 simply has to set her shares of those entries to 0. On the other hand, for entries M_{ij} of M such that features i and j are held by distinct parties, the parties need to compute an inner product between a column vector from X_1 and a column vector from X_2 . To do so, we rely on a secure inner product protocol which we present next (Section 3.1.1). We note also that because the two off-diagonal blocks of M are transpositions of

each other, computing only one of these blocks is enough to get an additive share of M .

In the multi-party case, similarly to the two-party case, the vertical partitioning of the database implies that each party holds a subset of the columns of X and a corresponding subset of the rows of X^\top . Since each value in A is an inner product between a row vector of X^\top and a column vector of X , this means that each value in A depends on the inputs of at most two parties. Thus, also in the multi-party case, the aggregation phase for vertically-partitioned datasets can be reduced to secure two-party computation of inner products. At the end of the aggregation phase, the parties will obtain A in a form where each entry is either known to one party or is shared between some pair of input parties.

3.1.1 Secure Inner Product

In this section, we present protocols for two parties P_1 and P_2 , holding vectors \vec{x} and \vec{y} , respectively, to securely compute the inner product of their vectors and obtain additive shares of the result. As mentioned before, we use a fixed-point encoding for real numbers (see Section 4 for details). Thus, our numbers can be represented as elements of a finite field \mathbb{Z}_q , and hence the functionality we need can be described as

$$f(\vec{x}, \vec{y}) = (r, \langle \vec{x}, \vec{y} \rangle - r) \quad (3)$$

where $\vec{x}, \vec{y} \in \mathbb{Z}_q^n$, r is a random element of \mathbb{Z}_q , and each party gets one of the components of the output. There are several techniques that can be used for this task as, essentially, it corresponds to secure multiplication.

Generally, MPC techniques represent the implemented functionality either as arithmetic circuits (e.g., SPDZ [15], TinyOT [58]) or as Boolean circuits (e.g., Yao’s protocol [70], GMW [31]). While the former makes it possible to represent arithmetic operations very efficiently, the latter is better for performing bit-level operations. A second important property of MPC protocols is round complexity. Most notably, some techniques such as Yao’s garbled circuits have a *constant* number of rounds, while for others (e.g., SPDZ, TinyOT, GMW), the number of rounds increases with the multiplicative depth of the circuit being evaluated (for Boolean circuits, this is the AND-depth). Note that the functionality in Equation (3) only consists of additions and multiplications. Furthermore, all multiplications can be done in parallel. Thus, a representation as an arithmetic circuit with constant multiplicative depth is the natural choice.

We propose two protocols for the functionality in Equation (3). The first is an extension of OT-based secure multiplication, which was originally proposed by Gilboa [29], and is used in different MPC contexts [18, 41]. It uses OT to compute

Parties: P_1, P_2 , and Trusted Initializer TI.

Inputs: $P_1 : \vec{x} \in \mathbb{Z}_q^n; P_2 : \vec{y} \in \mathbb{Z}_q^n$.

Outputs: $P_1 : s_1 \in \mathbb{Z}_q; P_2 : s_2 \in \mathbb{Z}_q$ such that $s_1 + s_2 = \langle \vec{x}, \vec{y} \rangle$.

Protocol:

1. TI generates random vectors $\vec{a}, \vec{b} \in \mathbb{Z}_q^n$ and a random number $r \in \mathbb{Z}_q$, and sets $z = \langle \vec{a}, \vec{b} \rangle - r$. It sends (\vec{a}, r) to P_1 , and (\vec{b}, z) to P_2 .
2. P_1 sends $\vec{x} + \vec{a}$ to P_2 .
3. P_2 sends $\vec{y} - \vec{b}$ to P_1 .
4. P_1 computes its output share as $s_1 = \langle \vec{x}, \vec{y} - \vec{b} \rangle - r$.
5. P_2 computes its output share as $s_2 = \langle \vec{x} + \vec{a}, \vec{b} \rangle - z$.

Fig. 1. Secure two-party inner product/multiplication supported by a trusted initializer. Its security relies on the fact that the TI does not collude with the data providers, as described in Lemma 1.

additive shares of the product of two numbers in a binary field, which can be trivially extended to securely compute the inner product of two vectors.

Our second secure inner product protocol uses only symmetric key operations and is much more efficient than the previous one. It builds upon techniques based on precomputation of multiplicative triples [2], as the secret sharing techniques mentioned above, and relies on an initializer that aids in the computation. Our protocol can also be seen as a modification of the construction presented by Du and Atallah [20] where the trusted initializer does not keep a share of the result, and is presented in Figure 1.

In both our protocols, we exploit the facts that, in our case, (i) the input vectors are each completely owned by one of the parties, and not shared among them, (ii) only two-party computations are needed, (iii) we are concerned about semi-honest security at this point. These observations lead to simpler protocols than general MPC protocols based on secret sharing.

In Section 6.3 we present an evaluation of our two protocols for secure inner product, in the context of our implementation of the aggregation phase for secure linear regression.

3.2 Solving Phase

After the aggregation phase, the parties hold additive shares of a square, symmetric, positive definite matrix A and a vector b , and the task is to securely solve $A\theta = b$.

In line with previous work [35, 59], and in order to achieve constant round-complexity, we choose Yao’s garbled circuits for our solving phase, and introduce two additional non-colluding parties. These aid in the computation without learning neither the result, nor anything about the parties’ in-

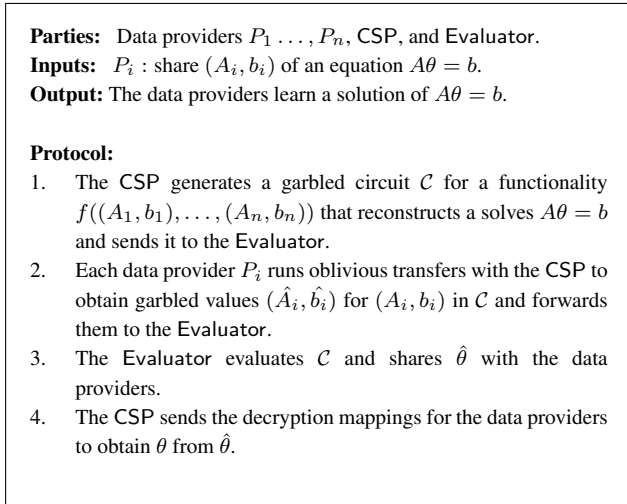


Fig. 2. The solving phase of our protocol. Its security in the semi-honest threat model follows from the security proof for Yao's protocol in [49].

put. In fact, our architecture relying on two external parties can be seen as a way of reducing a multi-party problem to a two-party problem, and had been used before in [59] and [56]. Analogously to the presentation in [59], we call our additional roles Crypto Service Provider (CSP) and Evaluator. For clarity, from now on, we will call our parties *data providers*. The CSP is in charge of generating a garbled circuit for solving $A\theta = b$, while the Evaluator will do the evaluation. In contrast to [59], the Evaluator does not necessarily learn the solution of the system in our protocol. The solving phase of our protocol is presented in Figure 2.

Note that the CSP and the Evaluator are only required to be non-colluding and hence, for the protocol in Figure 2, their roles could be taken by some of the data providers. In particular, for the case with two data providers, no additional parties are needed. We discuss security considerations of the solving phase protocol in the next section, where we present our complete protocol for secure linear regression.

Finally, let us note that we have not specified the details of the circuit \mathcal{C} yet. It is important to remark that, after having designed a fast aggregation phase, the bottleneck of our protocol in terms of running time will now be in the solving phase. Hence, the concrete algorithm for solving $A\theta = b$ is a crucial element for the scalability of our protocol for secure linear regression. In Sections 5 and 4 we discuss desirable properties of a good algorithm for solving systems of linear equations tailored for MPC, and propose a Conjugate Gradient Descent algorithm. Finally, let us mention that, as we will see in Section 5, the circuit \mathcal{C} has high degree. As mentioned in the previous section this motivates the choice of GCs as underlying

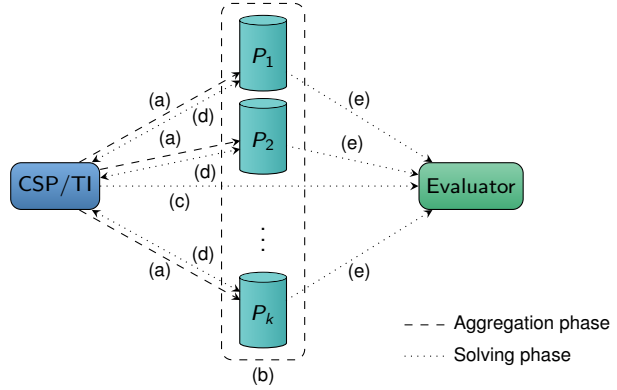


Fig. 3. The multi-party protocol: (a,b) The data providers compute additive shares of $A\theta = b$, by doing local precomputations and executing several instances of the inner product protocol of Figure 1. (c) The CSP generates a garbled circuit \mathcal{C} for a functionality that reconstructs and solves $A\theta = b$ and sends it to the Evaluator. (d) Each data provider P_i obtains garbled values for their shares using OT and forwards them to the Evaluator. (e) The Evaluator evaluates \mathcal{C} and shares the garbling of the result $\hat{\theta}$ with the data providers. Finally, The CSP sends the decryption mappings for the data providers to obtain θ from $\hat{\theta}$.

MPC technique. An extensive evaluation of MPC techniques for the task of linear system solving is left for further work.

3.3 Secure Linear Regression

Our main protocol is depicted in Figure 3. In this case, the composition between the aggregation and solving phases is implemented in steps (d) and (e) by means of oblivious transfers between the CSP and the data providers. Here, the roles of Trusted Initializer and CSP in the aggregation and solving phases are taken by a single additional non-colluding party, while the role of the Evaluator could be taken by one of the data providers.

In the two-party case, the roles of the CSP and the Evaluator can be taken by the data providers, and the composition of the two phases is straightforward, while the role of Trusted Initializer in the aggregation phase is taken by an additional non-colluding party that aids in the computation.

A variant of our protocol described above implements the aggregation phase using the protocol based on OT (Section 3.1.1). In this way, we remove the need for an external non-colluding party. We will consider the performance of this variant in the experimental evaluation, and revisit it when we consider extensions of our protocol to withstand malicious adversaries.

Up to this point we have focused on design and correctness aspects of our protocols, without discussing their security guarantees. We undertake that task in the following section.

3.3.1 Security Against Semi-Honest Adversaries

Our work is in the semi-honest – also known as *honest-but-curious* – threat model. Security in this model does not provide privacy, correctness, or even fairness guarantees if a party deviates from the protocol. However, the semi-honest setting is appropriate for distributed machine learning applications where parties are generally trusted, but should be prevented from sharing their data in the clear for legal reasons, or the computation is run in a controlled environment. This conditions are reasonably easy to meet in scenarios where the parties are corporations or government agencies. In this section we formally argue the correctness of our protocol in the semi-honest threat model, and in Section 7 we discuss an extension to guarantee correctness of the result.

Similarly to how we presented our protocols, we argue security of our linear regression protocol with a modular approach, by proving the security of the protocol based on the security of its subprotocols.

Let us thus first argue security of our protocols for the inner product functionality: while for the security of the OT-based protocol we refer the reader to [18, 29], we state the correctness and security of the protocol based on a Trusted Initializer in the following lemma.

Lemma 1. *The protocol in Figure 1 is a secure two-party computation protocol against semi-honest adversaries for the functionality $f(\vec{x}, \vec{y}) = (r, \langle \vec{x}, \vec{y} \rangle - r)$. where (\vec{x}, \vec{y}) and $(r, \langle \vec{x}, \vec{y} \rangle - r)$ are the inputs and outputs of the parties.*

Proof. Correctness can be easily verified from the definition of the resulting shares. To prove security, we construct simulators Sim_1 and Sim_2 , which simulate the views of the two parties.

Sim_1 generates a view for P_1 given \vec{x} and P_1 's output in the ideal functionality, denoted s_1^* . The message from the TI is simulated as $(\vec{a}', \langle \vec{x}, \vec{b}' \rangle - s_1^*)$, while the message from P_2 is simulated as \vec{b}' , where $\vec{a}', \vec{b}' \in U(\mathbb{Z}_q^n)$. The resulting view has the same distribution as P_1 's view in the real protocol execution.

Similarly, Sim_2 generates a view for P_2 given \vec{y} and s_2^* as follows: the message from the TI is simulated as $(\vec{b}', \langle \vec{a}', \vec{b}' \rangle - s_2^*)$, while the message from P_1 is simulated as \vec{a}' , where $\vec{a}', \vec{b}' \in U(\mathbb{Z}_q^n)$. These messages have the same distribution as P_2 's view in the real execution, which completes the proof. \square

The security of our main protocol depicted in Figure 3 is stated in the following theorem. Informally, the theorem states that, as long as the participants do not deviate from the protocol, and the external parties (CSP/TI and Evaluator) are non-colluding, none of the parties learn anything from the inputs

of the data providers that cannot be deduced from the result of the computation, namely the regression parameter θ . Moreover, recall that the external parties do not obtain θ .

Theorem 1 (Security). *Let Π_1 be a secure two-party computation protocol for the inner product functionality defined in Section 3.1.1 with security against semi-honest adversaries, and Π_2 be a two-party computation protocol based on garbled circuits with semi-honest security. The construction in Figure 3 is a secure computation protocol that provides security against semi-honest adversaries, in the setting where the CSP/TI and the Evaluator do not collude with neither the data providers nor each other.*

Proof. We need to show simulation security against a semi-honest non-colluding adversary that controls the CSP, a semi-honest non-colluding adversary that controls the Evaluator, and a semi-honest adversary that controls a subset of the input parties.

The security of the inner product protocol Π_1 implies that there exist simulators Sim_{P_1} and Sim_{P_2} which can simulate the view of each of the two parties in the execution of the protocol. The security of the two party computation protocol based on garbled circuits Π_2 implies the existence of simulators Sim_{garb} and Sim_{eval} which work as follows. Sim_{garb} simulates the view of the garbling party, which consists of the execution of the OT protocols as a sender. Sim_{eval} simulated the view of the Evaluator which consists of the garbled circuit itself as well as the view in the execution of the OT protocols as a receiver. We will use these simulators in our proof.

Semi-honest CSP. The view of the CSP in the secure computation protocol consists only of the messages exchanged in the garbled circuit computation where it has no input, and the messages exchanged in the OT executions with the data providers. Both can be simulated using Sim_{garb} .

Semi-honest Evaluator. Similarly to the case of the CSP, the view of the Evaluator in the execution of the protocol consists of the messages exchanged in the execution of the garbled circuit evaluation where it has no inputs. Therefore, we can use the simulator Sim_{eval} that can simulate its view.

Semi-honest Input Parties. The view of an adversary who controls a subset of the input parties in the protocol in Figure 2 consists of the messages exchanged in the executions of the inner product protocol in the aggregation phase with other input parties not controlled by the adversary, the messages in the OT executions with the CSP in the solving phase together with mappings for the output garbled labels as well as the garbled output the parties receive from the Evaluator. We can use the simulators Sim_{P_1} and Sim_{P_2} to obtain the view of the adversary in the aggregation phase. Additionally, we can use

Sim_{garb} for Π_2 to simulate the view of the adversary in the OT executions. Given the value θ of the output solution, which the parties will receive, the simulator Sim_{garb} can also generate mappings for the output garbled labels to real values and a set of garbled labels for the output that open to θ . This completes the proof. \square

The above security argument can be easily adapted to the case where the evaluator is one of the data providers, as long as CSP/TI does not collude with any other party.

We can further extend the security guarantees of our protocol, to the setting where the roles of the CSP and the Evaluator are executed by two of the parties, as long as these two parties are semi-honest and non-colluding with each other. In this scenario, we use the OT-based protocol for the aggregation phase, as it does not require a Trusted Initializer (which was instantiated by the CSP). The security guarantees for this variant in which *no external parties are required*, are stated in the next theorem.

Theorem 2 (Security without external parties). *Let Π be an instantiation of the protocol described in Figure 3 where the roles of CSP and the Evaluator are implemented by two distinct input providers, and the aggregation phase is instantiated with the OT-based inner product protocol. Then, Π provides security against semi-honest adversaries in the setting where the CSP and the Evaluator do not collude.*

In Section 6 we provide an experimental evaluation of the protocols of the two theorems above where we quantify the speed-up obtained by relying on an offline phase performed by a non-colluding external party acting as Trusted Initializer.

4 Number Representation

Implementing secure MPC protocols to work with numerical data poses a number of challenges. In our case, the most significant of those challenges is choosing an encoding for numerical data providing a good trade-off between efficiency and accuracy. The IEEE 754 floating-point representation is the standard choice used by virtually every numerical computation software because of its high accuracy and numerical stability when working with numbers across a wide range of orders of magnitude. However, efficient implementations of IEEE 754-compliant secure MPC protocols are a subject of on-going research [17, 61], with the current state-of-the-art yielding 32 bit floating-point multiplication circuits running in time comparable to computers from the 1960's [68]. Unfortunately, these implementations do not yet scale to the throughput required for data analysis tasks involving large datasets, and forces us

to rely on fixed-point encodings like previous work in this area [10, 35, 59]. The rest of this section presents the details of the particular fixed-point encoding used by our system, with a particular emphasis on its effect on the accuracy of the protocol. We further introduce data normalization and scaling steps that play an important role in the context of linear regression.

4.1 Fixed-Point Arithmetic

Our use of fixed-point encodings follows closely previous works on secure linear regression [10, 35, 59]. However, unlike these, we provide a formal analysis of the errors such encodings introduce when used to simulate operations on real numbers.

In a fixed-point signed binary encoding scheme each number is represented using a fixed number of bits b . These bits are subsequently split into three parts: one sign bit, b_f bits for the fractional part, and b_i bits for the integral part, with the obvious constraint $b = b_i + b_f + 1$. For each possible value of b_f and b_i one gets an encoding capable of representing all the numbers in the interval $[-2^{b_i} + 2^{-b_f-1}, 2^{b_i+1} - 2^{-b_f-1}]$ with accuracy 2^{-b_f} . In general, arithmetic operations with fixed-point numbers can be implemented using signed integer arithmetic, though special care is required to deal with operations that might result in overflow. We point the reader to [24] for an in-depth discussion of fixed-point arithmetic.

In order to implement the secure inner product protocol described in Figure 1 we need to simulate fixed-point arithmetic over a finite ring \mathbb{Z}_q whose elements can be represented using b bits. Thus, the encoding used by our protocols involves two steps: mapping reals into integers, and mapping integers to integers modulo q with $q = 2^b$. We introduce encoding and decoding maps for each of these two steps, as summarized in the following diagram:

$$\mathbb{R} \begin{array}{c} \xrightarrow{\phi_\delta} \\ \xleftarrow{\tilde{\phi}_\delta} \end{array} \mathbb{Z} \begin{array}{c} \xrightarrow{\varphi_q} \\ \xleftarrow{\tilde{\varphi}_q} \end{array} \mathbb{Z}_q .$$

The map encoding reals into integers is given by $\phi_\delta(r) = \lceil r/\delta \rceil$, where the parameter $\delta = 2^{-b_f}$ controls the encoding precision and $\lceil \cdot \rceil$ returns the rounding of a real number to the closest integer (with ties favouring the largest integer). In particular, δ is the smallest number that can be represented by our encoding ϕ_δ . The integer decoding mapping is given by $\tilde{\phi}_\delta(z) = z\delta$. Next we recall several well-known facts [24] about the error introduced by operating under a fixed-point encoding with finite precision δ . For any reals $r, r' \in \mathbb{R}$ we have the following:

1. $|r - \tilde{\phi}_\delta(\phi_\delta(r))| \leq \delta$,
2. $|(r + r') - \tilde{\phi}_\delta(\phi_\delta(r) + \phi_\delta(r'))| \leq 2\delta$,
3. $|(rr') - \tilde{\phi}_{\delta^2}(\phi_\delta(r)\phi_\delta(r'))| \leq (|r| + |r'|)\delta + \delta^2$.

Note the last fact involves a decoding with precision δ^2 to account for the increase in the number of fractional bits required to represent the product of two fixed-point numbers.

Encoding reals in a bounded interval $[-M, M]$ with ϕ_δ yields integers in the range $[-M/\delta] \leq \phi_\delta(r) \leq [M/\delta]$. This interval contains $K = 2M/\delta + 1$ integers. Thus, it is possible to map this interval injectively into any ring \mathbb{Z}_q of integers modulo q with $q \geq K$. We obtain such injection by defining the encoding map $\varphi_q(z) = z \bmod q$, which for integers in the range $-q/2 \leq z \leq q/2$ yields $\varphi_q(z) = z$ if $z \geq 0$ and $\varphi_q(z) = q + z$ for $z < 0$. The corresponding decoding map is given by $\tilde{\varphi}_q(u) = u$ if $0 \leq u \leq q/2$ and $\tilde{\varphi}_q(u) = u - q$ for $q/2 < u \leq q - 1$. Although φ_q is a ring homomorphism mapping operations in \mathbb{Z} into operations in \mathbb{Z}_q , decoding from \mathbb{Z}_q to \mathbb{Z} after operating on encoded integers might not yield the desired result due to overflows. To avoid such overflows one must check that the result falls in the interval where the coding φ_q is the inverse of $\tilde{\varphi}_q$. In particular, the following properties only hold for integers z, z' such that $|z|, |z'| \leq q/2$:

1. $\tilde{\varphi}_q(\varphi_q(z)) = z$,
2. $|z + z'| \leq q/2$ implies $z + z' = \tilde{\varphi}_q(\varphi_q(z) + \varphi_q(z'))$,
3. $|z \cdot z'| \leq q/2$ implies $z \cdot z' = \tilde{\varphi}_q(\varphi_q(z) \cdot \varphi_q(z'))$.

4.2 Accuracy of Inner Product

The properties of the encodings described in the previous section can be used to provide a bound on the accuracy of the result of the inner product protocol given in Figure 1. At the end of the protocol, both parties have an additive share of $\langle \tilde{x}, \tilde{y} \rangle$ for some integer vectors $\tilde{x}, \tilde{y} \in \mathbb{Z}_q^n$. Therefore, we want to show that when the input vectors \tilde{x}, \tilde{y} provided by the parties are fixed-point encodings of some real vectors $\vec{u}, \vec{v} \in \mathbb{R}^n$, then the result of the inner product over \mathbb{Z}_q can be decoded into a real number approximating $\langle \vec{u}, \vec{v} \rangle$. We note that the comparison here is against real arithmetic, while in the experimental section we will be comparing against floating-point arithmetic, which can represent real arithmetic to high degrees of accuracy across a much larger range of numbers [43].

According to the previous section, there are two sources of error when working with fixed-point encoded real numbers: the systematic error induced by working with finite precision (which can be mitigated by increasing b_f), and the occasional error that occurs when the numbers in \mathbb{Z}_q overflow (which can be mitigated by increasing b_i). To control the overflow error we assume that \vec{u} and \vec{v} are n -dimensional real vectors with entries bounded by R , $|u_i|, |v_i| \leq R$ for $i \in [n]$, and $\tilde{x} = \varphi_q(\phi_\delta(\vec{u}))$ and $\tilde{y} = \varphi_q(\phi_\delta(\vec{v}))$ are encodings of those vectors with precision δ .

Theorem 3 ([69, Ch. 3]). *If $q > 2nR^2/\delta^2$, then:*

$$|\langle u, v \rangle - \tilde{\phi}_{\delta^2}(\tilde{\varphi}_q(\langle \tilde{x}, \tilde{y} \rangle))| \leq 2nR\delta + n\delta^2 .$$

The proof is straightforward and is provided in Appendix A for the sake of completeness. We note that the assumption of a bound R on the entries of the vectors \vec{u} and \vec{v} is not very stringent: if both parties agree on a common bound R , or there is a publicly known R , then the vectors can be normalized locally by each party before the execution of the protocol. Thus, in practice it is convenient to normalize the real vectors u and v such that their entries are bounded by C/\sqrt{n} for some constant $C > 0$. In this case, the bounds above can be used to show that if one requires the final inner product to have error at most ε , this can be achieved by taking $\delta = \varepsilon/2C\sqrt{n}$ and $q = 8C^2n/\varepsilon^2$. Using these expressions we see that an encoding with $b = O(\log(n/\varepsilon))$ bits is enough to achieve accuracy ε when performing n -dimensional inner products with fixed-point encoded reals.

4.3 Data Standardization and Scaling

The statistical theory behind ridge regression generally assumes that the covariates in the columns of X are normally distributed with zero mean and unit variance. In practical applications this is not always satisfied. However, it is a common practice to standardize each column of the training data to have zero mean and unit variance. This column-wise procedure is implemented as part of the local pre-processing done by each party in our linear regression protocol. We use \bar{X} to denote the data matrix obtained after standardization.

The importance of proper data scaling to prevent overflows in the aggregation phase is the key message from Theorem 3 (Section 4.2). Additionally, a critical phenomenon occurring in the solving phase is the degradation in the quality of the solution when the encoding is fixed and the dimension of the linear system grows. To alleviate the effect of the dimension on the accuracy of the solving phase, our pre-processing re-scales every entry in \bar{X} and Y by $\frac{1}{\sqrt{d}}$. We note that these scalings can be performed locally by the parties holding each respective column.

The standardization and scaling steps described above imply that our protocols effectively solve the linear system

$$\left(\frac{1}{nd} \bar{X}^\top \bar{X} + \lambda I \right) \theta = \frac{1}{nd} \bar{X}^\top Y .$$

This scaling has no effect on the linear regression problem being solved beyond changing the scale of the regularization parameter. On the other hand, the standardization step needs to be reversed if one wants to make predictions on test data following the same distribution as the original training data. This

task corresponds to a simple linear computation that can be efficiently implemented in MPC, both in the case where the model θ is disclosed to the parties and the case where it is kept shared among the parties and the prediction task using θ is itself implemented in a secure way using MPC.

5 Solving Linear Systems

As discussed in Section 3.2, the solving phase of our protocol involves solving positive definite linear systems of the form $A\theta = b$ in a garbled circuits protocol. A wide variety of solutions to this problem can be found in the numerical analysis literature, though not all of them are suitable for MPC. For example, any variant of Gaussian elimination involving data-dependent pivoting can be immediately ruled out because implementing non-data-oblivious algorithms inside garbled circuits produces unnecessary blow-ups in the circuit size. This limitation has already been recognized by previous works on private linear regression using garbled circuits and other MPC techniques [35, 59]. In particular, these works consider two main alternatives: computing the full inverse of A , or solving the linear system directly using the Cholesky decomposition of A as an intermediate step. Although both of these methods have an asymptotic computational cost of $O(d^3)$, in practice Cholesky is faster and numerically more stable. This justifies the choice of Cholesky as a procedure for solving linear systems in garbled circuits made in [59]. However, as we show in Section 6.2, Cholesky’s cubic cost in the dimension can become prohibitive when working with high-dimensional data, in which case one must resort to iterative approximate algorithms like Conjugate Gradient Descent (CGD).

In the following Section we present a novel CGD algorithm tailored for working with fixed-point arithmetic and compare it to the classical CGD algorithm.

5.1 Conjugate Gradient Descent

Factorizing the coefficient matrix A in order to make the solution easier to compute, as in Cholesky’s algorithm, is not the only efficient way to solve a system of linear equations. An entirely different approach relies on iterative algorithms which construct a monotonically improving sequence of approximate solutions θ_t converging to the desired solution. The main virtue of iterative algorithms is the possibility to reduce the cost of solving the system below the cubic complexity required by exact algorithms at the expense of providing only an approximate solution. In particular, for linear systems arising from ridge regression the practical importance of iterative




Fig. 4. Fixed-point CGD Implementations (with 64 bits on a system of dimension 50). (left) Textbook CGD [60]. (right) Our Fixed-point CGD.

algorithms is twofold: since intensive computations inside an MPC framework can be expensive, iterative methods provide a natural way to spend a fixed computational budget on finding an approximation to the desired solution by stopping after a fixed number of iterations; and, since the coefficients of the linear system occurring in ridge regression are inherently noisy because they are computed from a finite amount of data, finding an approximate solution whose accuracy is on the same order of magnitude as the noise in the coefficients is generally sufficient for optimizing the predictive performance of the model.

When the coefficient matrix of the linear system $A\theta = b$ is positive definite, a popular iterative method is the *conjugate gradient descent* (CGD) algorithm. The CGD algorithm can be interpreted as solving the system by iteratively minimizing the objective $\|A\theta - b\|^2$ with respect to the parameter vector θ using the method of conjugate gradients [60]. However, the numerical stability of CGD is in general worse than that of Cholesky’s algorithm, making it very sensitive to the choice of fixed-point encoding. This can be observed in Figure 4 (left), where we plot the behavior of the residual $\|A\theta_t - b\|$ as a function of t for several settings of the number of integer bits b_i when solving a system with $d = 50$ using 64-bit fixed-point encodings. This plot shows that CGD is very sensitive to the choice of b_i , and even in the optimal setting ($b_i = 10$) the convergence rate is much slower than the one achieved by the corresponding floating-point implementation (dashed black curve).

After a thorough investigation of the behavior of a fixed-point implementation of CGD we concluded that the main problem occurs when the norm of the conjugate gradients decreases too fast and consequently the step sizes grow at each iteration. This motivates a variant of CGD, which we call *fixed-point CGD* (FP-CGD), and whose pseudo-code is given in Figure 5. The only difference between standard CGD and FP-CGD is the use of normalized gradients $\tilde{g}_t = g_t / \|g_t\|_\infty$ in the computation of the conjugate search directions p_t ; in particular, by taking $\tilde{g}_t = g_t$ one recovers the classical textbook

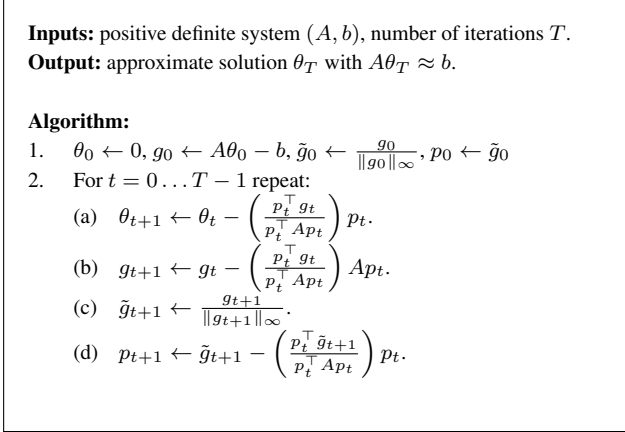


Fig. 5. Our variant of the Conjugate Gradient Descent algorithm, which is optimized for the use with fixed-point encoded numbers as described in Section 5.1.

algorithm. While CGD has been used in the context of secure computation in previous work [6], to the best of our knowledge, this modification of CGD has not been considered in the literature before.

It is easy to show that CGD and FP-CGD produce exactly the same sequence of approximate solutions θ_t when working in exact arithmetic. However, normalizing the search directions in FP-CGD makes the method much more stable with respect to the errors introduced by fixed-point arithmetic. This is illustrated by Figure 4 (right), where we show the evolution of the residuals for the same system and fixed-point encodings that we used to test standard CGD. Here, we observe that we recover the same converge rate as CGD with floating-point arithmetic while only suffering a loss between 2 and 3 digits of accuracy. We will see later in our experiments with real data (Section 6.4) that this loss in accuracy is negligible in real applications, and it is comparable (or better) than the accuracy provided by a fixed-point implementation of Cholesky’s method. In terms of computation, we note that each iteration of FP-CGD is slightly more expensive than an iteration of standard CGD due to the normalization step, but the asymptotic cost per iteration is in $\Theta(d^2)$ in both cases.

Further justification for the use of FP-CGD instead of other iterative methods is provided by its favorable theoretical properties. In exact arithmetic, FP-CGD inherits two important properties from CGD: (i) it converges to the *exact solution* of $A\theta = b$ in exactly d iterations; and (ii) it achieves the *optimal convergence rate* among all first-order methods, yielding a solution with error at most ε after $O(\sqrt{\kappa(A)} \log(1/\varepsilon))$ iterations. Property (i) is important because not all iterative algorithms for solving linear systems are guaranteed to produce an exact solution after a finite number of iterations, including

standard gradient descent methods like the ones implemented in [33] using leveled homomorphic encryption. Furthermore, property (ii) says essentially that it is not possible to find a first-order method producing more accurate approximate solutions than CGD, and that only a few iterations of CGD are required to accurately solve linear systems with small condition number $\kappa(A)$. In principle, these properties might not be preserved by the fixed-point implementations of CGD and FP-CGD due to the numerical errors introduced by the finite-precision arithmetic. In the case of CGD, [52] shows that when working with a finite number of bits for the fractional part and an infinite number of bits for the integer part, several important properties of the exact version, including the convergence rate, are preserved up to small order modifications. However, the assumption of infinite bits for the integer part is not realistic in practice, and the experiments in Figure 4 show that when a fixed number of bits needs to be split between the integer and fractional part of the representation a crucial tension arises. FP-CGD is designed to alleviate such tension by making sure that a minimal number of bits in the integer part suffice to solve a properly scaled system accurately. We demonstrate this through our experimental evaluations in Section 6, where we also compare the accuracy and running time of FP-CGD against the method based in Cholesky decomposition used in [59], and study the effect of the condition number $\kappa(A)$ on both methods.

6 Experimental Results

This section presents an extensive empirical evaluation of our protocols. We start by describing the details about the concrete implementation of our system and the experimental setup. Then we present experiments for different implementations of the solving phase using FP-CGD and Cholesky, comparing them in terms of accuracy and running time across a wide range of settings. The second set of experiments evaluates two implementations for the aggregation phase: one based on the inner product protocol relying on a trusted initializer, and the one based on OT. Finally, we present an evaluation of the complete system on nine real datasets from the UCI repository with different algorithms in the solving phase and compare the resulting predictive performance of the learned model against the performance obtained in the non-private setting. Overall, our experiments highlight the importance of using FP-CGD for high-dimensional data, and provide a guide on how to choose a good fixed-point encoding depending on the characteristics of the problem.

6.1 Implementation and Setup

We implemented our MPC protocols using Obliv-C [72], an extension of the C programming language for secure computation. Obliv-C includes an implementation of Yao’s garbled circuit protocol that incorporates recent optimizations including free XOR [44], garbled row reduction [56], fixed key block ciphers [3], and half gates [73]. Further, it features efficient implementations of OT extension, including variants such as correlated OT [1].

To support arbitrary precision arithmetic, we implemented a big integer library [74] for multi-party computation as an extension of Obliv-C. This library composes an arbitrary number of garbled integers to represent larger integer values in two’s complement binary. It features standard operations such as comparisons, bit manipulation and shifts, and arithmetic functions including addition, multiplication, integer division and modulo computation. All the arithmetic operations are implemented using common, efficient algorithms for extended precision arithmetic; in particular, we implement the Karatsuba-Comba [39] method for multiplication, and Knuth’s algorithm D [43] for division.

For fixed-point arithmetic, we implement two variants of our protocols. One uses Obliv-C’s native integer types to represent fixed-point numbers. Since for fixed-point multiplication, the bit width of intermediate values can be large as the sum of the arguments’ bit widths, this version is limited to 32-bit numbers. Our second fixed-point implementation uses the big integer library described above, allowing to represent numbers with arbitrary precision. This arrangement incurs some overhead relative to using Obliv-C’s native types. However, this overhead is justified in cases where 32-bit arithmetic introduces large errors. Although in principle, our second implementation allows for an arbitrary number of bits, we only evaluate it with 64 bit numbers. As the results in Section 6.4 show, this is enough to get accurate results on real datasets.

For the solving phase, we implemented the Cholesky and FP-CGD methods described by Nikolaenko et al. [59] and Section 5, respectively. Cholesky decomposition requires square root computation. Nikolaenko et al. use an iterative algorithm for computing square roots in garbled circuits based on Newton’s method. Our implementation of Cholesky’s method follows closely their approach and is based on the pseudo-code given in [59].

For the experiments in Sections 6.2 and 6.3, where each phase of the protocol is evaluated separately, we use synthetically generated data. For each setting of d and n , we sample n data points $x^{(i)}$ from a standard d -dimensional Gaussian distribution and a d -dimensional vector of parameters θ^* with independent coordinates sampled uniformly in the interval $[0, 1]$. The training labels are obtained as $y^{(i)} = \langle \theta^*, x^{(i)} \rangle + \epsilon^{(i)}$,

where $\epsilon^{(i)}$ is a noise term sampled from a Gaussian distribution with zero mean and variance $\sigma^2 = 0.1$. The regularization parameter in the solving phase is set to $\lambda = \sigma^2 d/n \|\theta^*\|^2$, which is the optimal choice suggested by the statistical theory of ridge regression.

For the experiments with synthetic data using 64 and 32 bit fixed-point encodings we used 60 and 28 bits for the fractional part, respectively. In the experiments with real data, the split between the fractional and integer part of the fixed-point encoding was optimized individually for each problem to obtain the best predictive performance.

The experiments were executed using Amazon EC2 C4 instances, each having 4 CPU cores, 7.5 GiB of RAM and 1 Gbps bandwidth. All our code, including the programs used for synthetic data generation, is publicly available under open source licenses [75].

6.2 Solving Phase

Figure 6 (left) compares the running times of the solving phase using Cholesky and FP-CGD for dimensions ranging from 10 to 500. For FP-CGD, we give the running time for 5, 10, 15, and 20 iterations. The top plot corresponds to the 64-bit implementation and the bottom plot to the 32-bit implementation. While Cholesky is faster than CGD for lower values of d , its cubic dependence on the dimension makes running a fixed number of iterations of FP-CGD faster as d increases. Thus, for high-dimensional data the iterative FP-CGD method is preferable in terms of computation time. The time spent running oblivious transfers is also shown, and accounts for a small fraction of the running time. For example, for the 64-bit implementations with $d = 200$, FP-CGD with 15 iterations runs in less than 45 minutes, while Cholesky takes more than an hour and a half. For $d = 500$, Cholesky takes more than 24 hours while FP-CGD with 15 iterations takes less than 4.5 hours. Similar results reporting circuit size instead of running time are presented in Appendix B due to space limitations (Figure 7).

Next, we compare both algorithms in terms of accuracy of the results. Here, the error is measured using the Euclidean distance to the optimal solution, obtained via floating-point computation. Figure 6 (middle) shows how the accuracy of the result is affected by the condition number of the input matrix A , comparing Cholesky with FP-CGD for varying numbers of iterations on problems with $d = 20$. We observed that FP-CGD achieves the same accuracy as Cholesky in the 64-bit version, and has even lower error when using only 32 bits. Moreover, the influence of the input condition number decreases with the number of iterations. After convergence, the error remains the same for all tested condition numbers. This advantage over

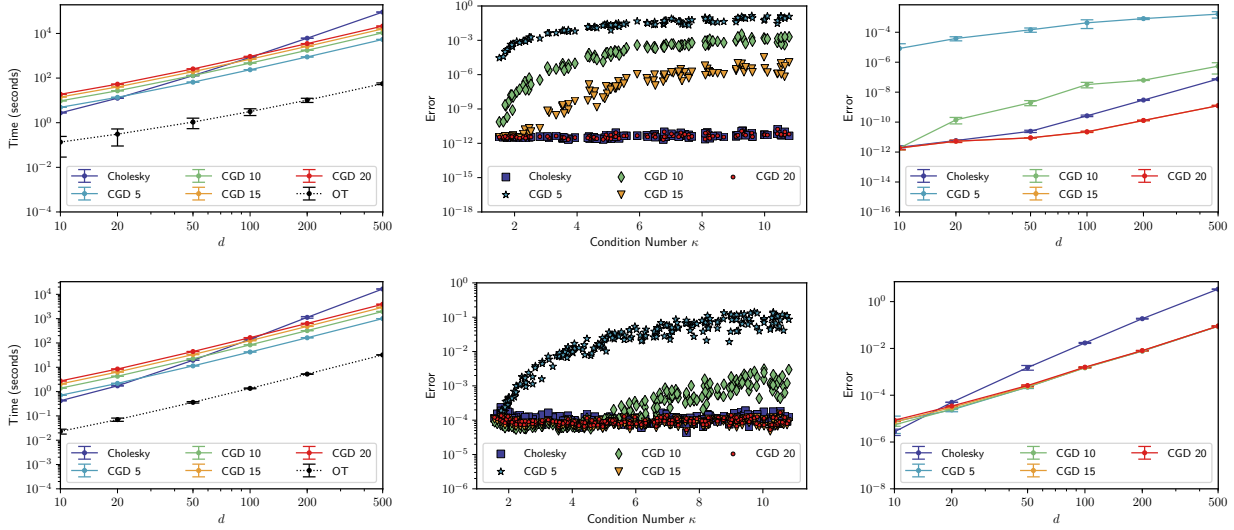


Fig. 6. (Left) Comparison between different methods for solving linear systems: Running time (seconds) of Cholesky and FP-CGD (with 5, 10, 15, and 20 iterations) as a function of input dimension. (Middle) Accuracy of Cholesky and CGD depending on the condition number of the input matrix A with $d = 20$. (Right) Accuracy of Cholesky and CGD, as a function of the input dimensionality d . (Top) Fixed-point numbers with $b = 64$ bits, with $b_f = 60$ in the fractional part. (Bottom) $b = 32$, $b_f = 28$.

Cholesky increases with higher d , as is shown in Figure 6 (right). For both the 32-bit and 64-bit versions, FP-CGD is more accurate than Cholesky as soon as $d > 50$.

6.3 Aggregation Phase

Table 1 shows a comparison between the running times of the 64-bit versions of our two aggregation protocols: *OT* is the protocol using the OT-based inner product protocol, and *TI* is the protocol that leverages a Trusted Initializer for the inner product protocol. We vary the number of records, n , from 50,000 up to one million, and the number of features, d , from 20 to 200. As expected, the protocol that takes advantage of the Trusted Initializer performs significantly better in all cases. However, since the TI's resources are the bottleneck in this setting, the protocol does not scale well to multiple data providers, as the fraction of the coefficient matrix A that can be computed locally shrinks as the number of parties increases. The OT-based version, on the other hand, handles many parties very well, due to the fact that OTs between different pairs of parties can be performed in parallel. Timing results our protocols for the aggregation phase for a more extensive set of configurations can be found in Tables 4 and 5 of Appendix B.

As a baseline, we implemented a garbled circuit protocol for performing a *single* inner product in Obliv-C. The running time for $n = 10^6$ was over 90 minutes. This implies that our protocol outperforms this naïve approach by several orders of magnitude.

n	d	Number of parties					
		2		3		5	
		OT	TI	OT	TI	OT	TI
$5 \cdot 10^4$	20	1m50s	1s	1m32s	2s	1m7s	2s
$5 \cdot 10^4$	100	42m12s	25s	34m39s	32s	24m58s	37s
$5 \cdot 10^5$	20	18m18s	15s	14m29s	18s	12m10s	21s
$5 \cdot 10^5$	100	7h3m56s	4m47s	5h20m52s	6m1s	4h17m8s	6m58s
$1 \cdot 10^6$	100	-	10m1s	-	12m42s	-	14m48s
$1 \cdot 10^6$	200	-	39m16s	-	49m56s	-	59m22s

Table 1. Comparison of running times between OT-based (left) and TI-based (right) aggregation protocols using 64 bit numbers. The running time of the Trusted Initializer, which is an offline pre-processing phase, is included. The complete results with additional parameter values can be found in Appendix B.

6.4 Experiments on Real Datasets

Although we have discussed the accuracy of our aggregation and solving protocols independently, we still have to evaluate our protocol in the task of building a ridge regression model. We evaluate our secure multi-party ridge regression system on 9 different regression problems selected from the UCI repository [46]. Each problem comes with a set of n examples of some dimension d which are randomly split into training (70%) and test sets (30%). Details about the names, original references where the dataset appeared, dimensions and number of examples in each task are given in Table 2. The dimensions of the problems range from 7 to 384, and the number of training examples ranges from over 200 to almost 3 million.

Examples in the test set are used to evaluate the predictive accuracy of the learned models in terms of their root mean squared error (RMSE). The predictive accuracy of the model will depend on the particular regularization parameter λ used in the regression: larger values prevent overfitting and should be used on small datasets, while smaller values reduce the bias in the regression and should be used on large datasets. Hyper-parameter tuning algorithms can use a set of test examples to evaluate the predictive power obtained with different regularization parameters and find a near-optimal setting. In principle, any hyper-parameter tuning algorithm based on solving the regression multiple times with different regularization parameters could be run on top of our system at the cost of executing the protocol repeatedly. This would involve many repeated computations but it is also possible to design a variation of our protocol to avoid most of these repetitions, in particular the computation of $X^\top X$. We leave this optimization for future work and concentrate on evaluating our ridge regression protocol with a fixed regularization λ for each problem selected a priori using the given train and test split. Table 2 gives the values selected for each task in addition to the condition number of the resulting linear system.

In each case, the features were split evenly between two parties. The scalings and normalizations described in Section 4.3 are performed locally by each party during the aggregation phase. We use the version of our protocol with the aggregation phase implemented using the inner product protocol of Figure 1, and the reported bit width of 32 or 64 bits in each case. We compare four implementations for the solving phase: Cholesky and Fixed-Point CGD, both in their 32 and 64 bits versions. In the case of CGD, we fixed the number of iterations to 20, regardless of the dimension or any other feature of the problem. For each task, we select the number of bits b_f for the fractional part in order to make sure we have enough bits b_i in the integer part for representing the largest coefficient appearing in the linear system (A, b) . The particular settings for each problem are also given in Table 2.

The results are presented in Table 3. We give the total running time of the protocol and RMSE on the test set for each of the four different implementations of the solving phase. The RMSE of each algorithm is compared to the one obtained using an insecure ridge regression algorithm implemented in Matlab that uses 64-bit floating-point representations (given in the column headed *Optimal*). The percentages in parentheses next to each RMSE give the relative increase – or in very few cases, decrease – on the error incurred using the private algorithm. Remarkably, the differences in accuracy are almost negligible for both 64 bit implementations, while in the 32 bit implementations, CGD is better than Cholesky except in one task.

In terms of running time, we note that both bit settings observe the same pattern: Cholesky is faster for small dimensions $d \leq 100$, and CGD is faster for large dimensions $d > 100$. However, even in the cases where Cholesky is faster, CGD is not far behind (a bit over 3 minutes in the worst case). On the other hand, in the cases where CGD is faster Cholesky has prohibitive running times, e.g., CGD takes ~ 4 h on the largest problem ($d = 384$) while Cholesky takes practically half a day.

Overall, we observe that for except for the last task, Fixed-point CGD with 32 bits always obtains relative errors below 1% and runs in less than 45 minutes, thus offering a very competitive solution. For the last dataset, that has large number of training examples, 64 bits are required to achieve good accuracies, mainly due to the errors introduced in the aggregation phase when the number of examples is large and there are not enough bits to represent the intermediate values of the computation of $X^\top X$ with sufficient precision. A possible workaround to this problem would require having a different number of bits in the aggregation and solving phases; we will consider this optimization in future work.

7 Beyond Semi-Honest Security

The protocols that we discussed in Section 3 guarantee security in an adversarial setting where all parties, including the data providers, the CSP, and the Evaluator, are semi-honest, and the CSP and the Evaluator do not collude. While security against semi-honest adversaries may be enough in a federated learning setting as argued in Section 3.3.1, some applications require stronger guarantees.

In this section, we propose an extension of our protocol that allows the data providers to check the correctness of the result of the second phase, which is computed by CSP and Evaluator. This correctness check works even if one of these two is malicious, or *actively* corrupted, i.e., they may deviate arbitrarily from the protocol description. We stress, however, that this correctness check is *not* enough to satisfy the definition of security against malicious adversaries as given by Goldreich [30]. In particular, a malicious CSP can cause the correctness check to succeed or fail depending only on one of the parties' private inputs. This *selective failure* attack has been discussed thoroughly in the context of garbled circuits [38, 53].

Our correctness check comes in the form of a third phase, which we call the *verification phase*. It is implemented using another small garbled circuit, which is secure against malicious adversaries. The general idea of using a tailored, lightweight verification procedure with high security guaran-

id	Name	Reference	d	n	λ	$\kappa(A)$	$b_f(b = 32)$	$b_f(b = 64)$
1	Student Performance	[11, 14]	30	395	$1.4 \cdot 10^{-2}$	$5.5 \cdot 10^0$	30	62
2	Auto MPG	[76]	7	398	$2.2 \cdot 10^{-3}$	$9.0 \cdot 10^1$	28	61
3	Communities and Crime	[63, 64]	122	1994	$2.0 \cdot 10^{-4}$	$1.1 \cdot 10^3$	24	53
4	Wine Quality	[12, 13]	11	4898	$2.9 \cdot 10^{-3}$	$6.8 \cdot 10^1$	29	60
5	Bike Sharing Dataset	[25, 26]	12	17 379	$8.2 \cdot 10^{-7}$	$2.2 \cdot 10^2$	28	60
6	Blog Feedback	[8, 9]	280	52 397	$1.1 \cdot 10^{-5}$	$1.3 \cdot 10^4$	25	54
7	CT slices	[34]	384	53 500	$9.3 \cdot 10^{-6}$	$1.6 \cdot 10^4$	25	51
8	Year Prediction MSD	[4]	90	515 345	$1.1 \cdot 10^{-5}$	$1.7 \cdot 10^2$	26	58
9	Gas sensor array	[27, 28]	16	4 208 261	$4.1 \cdot 10^{-7}$	$1.2 \cdot 10^5$	26	53

Table 2. Specifications of UCI datasets considered in our evaluation. The number of samples n is split randomly into training (70%) and test sets (30%). For each dataset, the regularization parameter λ and the number of fractional bits b_f were chosen as described in Section 6.1. The condition number κ was computed after data standardization and scaling (Section 4.3).

id	Optimal		FP-CGD (32 bits)		Cholesky (32 bits)		FP-CGD (64 bits)		Cholesky (64 bits)	
	RMSE	time	RMSE	time	RMSE	time	RMSE	time	RMSE	time
1	4.65	19s	4.65 (-0.0%)	5s	4.65 (-0.0%)	1m53s	4.65 (-0.0%)	35s	4.65 (-0.0%)	
2	3.45	2s	3.45 (-0.0%)	0s	3.45 (-0.0%)	13s	3.45 (0.0%)	1s	3.45 (0.0%)	
3	0.14	4m27s	0.14 (0.3%)	4m35s	0.14 (-0.0%)	24m24s	0.14 (0.2%)	26m31s	0.14 (-0.0%)	
4	0.76	3s	0.76 (-0.0%)	0s	0.80 (4.2%)	23s	0.76 (-0.0%)	4s	0.76 (-0.0%)	
5	145.06	4s	145.07 (0.0%)	1s	145.07 (0.0%)	26s	145.06 (0.0%)	4s	145.06 (0.0%)	
6	31.89	24m5s	31.90 (0.0%)	53m24s	32.19 (0.9%)	2h3m39s	31.90 (0.0%)	4h40m23s	31.89 (-0.0%)	
7	8.31	44m46s	8.34 (0.4%)	2h13m31s	8.87 (6.7%)	3h51m51s	8.32 (0.1%)	11h49m40s	8.31 (-0.0%)	
8	9.56	4m16s	9.56 (0.0%)	3m50s	9.56 (0.0%)	16m43s	9.56 (0.0%)	13m28s	9.56 (0.0%)	
9	90.33	48s	95.05 (5.2%)	42s	95.06 (5.2%)	1m41s	90.35 (0.0%)	1m9s	90.35 (0.0%)	

Table 3. Results of the evaluation of our system on UCI datasets. For each choice of algorithm and bit width, running time is reported, and the root mean squared error (RMSE) of the solution obtained by our system and an insecure implementation of ridge regression using floating point are compared.

tees to ensure correctness of a protocol with semi-honest security has been used by de Hoogh et al. [16] for linear programming, Laud and Pettai [45] for sorting, and Nikolaenko et al. [59] for ridge regression.

7.1 The Verification Phase

First, note that for the aggregation phase, we need to use the inner product protocol based on OT (Section 3.1.1), since we can no longer assume that the CSP can act as a trusted initializer. Since neither the CSP nor Evaluator take part in this aggregation phase, it needs no further analysis.

Nikolaenko et al. [59] made a perceptive observation that in the setting of linear regression, the correctness of the result can be verified simply by checking that the solution θ of the system $A\theta = b$ indeed minimizes the least squares expression of equation (1). This is done by verifying that θ evaluates to 0 in the derivative of the least squares function in equation (2), which is much cheaper than computing θ with malicious security. Concretely, this means checking that $\|2(A\theta - b)\| = 0$.

As we are working with finite-precision arithmetic, the equality check must be approximated as $\|2(A\theta - b)\| \in [-u, u]$, for some u chosen by the parties. We assume that the election of u is done correctly, in the sense that, if the CSP does not deviate from the protocol, then $\|2(A\theta - b)\| \in [-u, u]$ holds. If we consider the infinity norm, then the verification check of a solution θ , corresponds to checking

$$\forall i \in [d] : \vec{v}_i \in [-u, u] \quad (4)$$

where $\vec{v} = A\theta - b$, and A, b are additively shared among the data providers as $(A_1, b_1), \dots, (A_k, b_k)$ as a result of the aggregation phase. We call this check the *verification phase*, which is run after the solving phase. Hence, as mentioned above, we use a semi-honest protocol for the solving phase, and then run a verification protocol with malicious security.

In our protocol for the semi-honest case of Figure 3, θ is revealed to the data providers as a result of the solving phase. Consequently, an additive share of \vec{v} in (4) can be precomputed locally by the parties, and hence securely evaluating (4) only requires additions and comparisons. Using this, we implemented the verification phase as follows.

1. The data providers generate a uniformly pseudorandom vector $\vec{v}_{i,1} \in \mathbb{Z}_q^d$ and locally compute $\vec{v}_{i,2} = A_i\theta - b_i - \vec{v}_{i,1}$. Then, they send $\vec{v}_{i,1}$ to the CSP and $\vec{v}_{i,2}$ to the Evaluator.
2. CSP and Evaluator add up their received shares and obtain $\vec{v}_1 = \sum_{i=1}^k \vec{v}_{i,1}$ and $\vec{v}_2 = \sum_{i=1}^k \vec{v}_{i,2}$, and then run a two-party garbled circuit protocol with malicious security. In the circuit, $\vec{v} = \vec{v}_1 + \vec{v}_2$ is recovered, and a single bit is returned, indicating whether (4) holds.
3. CSP and Evaluator send the result of the verification circuit to all data providers.

Note that since not both CSP and Evaluator cannot be malicious at the same time, the parties only need to check if both bits they received are 1.

We implemented and evaluated the garbled circuit with malicious security for this verification phase, as an extension for our solving phase, using the EMP framework [67]. As expected, it runs extremely fast: in our experiments, garbling and execution took less than 3 seconds for $d \leq 500$. This is in contrast with the time needed for the solving phase for $d = 500$: 30 minutes, for 10 iterations of CGD and a bit width of 32 bits (see Figure 6, left).

8 Discussion

The problem of securely running machine learning algorithms when the training data is distributed among several parties is an important milestone for the development of privacy-preserving data analysis tools. In this paper, we focus on a linear regression task widely used in practical applications. Beyond the settings described in this paper, our implementation of secure conjugate gradient descent with fixed-point arithmetic and early stopping can also be used to deal with non-linear regression problems based on kernel ridge regression, given MPC protocols for evaluating kernel functions typically used in machine learning. Moreover, as mentioned in Section 3, an extensive evaluation of MPC techniques for the task of linear system solving, including our conjugate gradient descent algorithm, is an interesting continuation of the work proposed here. From a more theoretical perspective, the problem of providing security guarantees against malicious adversaries for approximate MPC functionalities poses interesting open challenges both in general and from the perspective of concrete machine learning tasks.

References

- [1] G. Asharov, Y. Lindell, T. Schneider, and M. Zohner. More efficient oblivious transfer and extensions for faster secure computation. In *ACM Conference on Computer and Communications Security*, pages 535–548. ACM, 2013.
- [2] D. Beaver. Efficient multiparty protocols using circuit randomization. In *CRYPTO*, volume 576 of *Lecture Notes in Computer Science*, pages 420–432. Springer, 1991.
- [3] M. Bellare, V. T. Hoang, S. Keelveedhi, and P. Rogaway. Efficient garbling from a fixed-key blockcipher. In *IEEE Symposium on Security and Privacy*, pages 478–492. IEEE Computer Society, 2013.
- [4] T. Bertin-Mahieux. YearPredictionMSD data set. <https://archive.ics.uci.edu/ml/datasets/YearPredictionMSD>, 2011.
- [5] M. Blanton, A. Steele, and M. Aliasgari. Data-oblivious graph algorithms for secure computation and outsourcing. In *ASIACCS*, pages 207–218. ACM, 2013.
- [6] D. Bogdanov, L. Kamm, S. Laur, and V. Sokk. Rmind: a tool for cryptographically secure statistical analysis. *IEEE Transactions on Dependable and Secure Computing*, 2016.
- [7] D. Bogdanov, S. Laur, and J. Willemson. Sharemind: A framework for fast privacy-preserving computations. In *ESORICS*, pages 192–206. Springer, 2008.
- [8] K. Buza. Feedback prediction for blogs. In *GfKI, Studies in Classification, Data Analysis, and Knowledge Organization*, pages 145–152. Springer, 2012.
- [9] K. Buza. BlogFeedback data set. <https://archive.ics.uci.edu/ml/datasets/BlogFeedback>, 2014.
- [10] M. D. Cock, R. Dowsley, A. C. A. Nascimento, and S. C. Newman. Fast, privacy preserving linear regression over distributed datasets based on pre-distributed data. In *AI@Sec@CCS*, pages 3–14. ACM, 2015.
- [11] P. Cortez. Student performance data set. <https://archive.ics.uci.edu/ml/datasets/Student+Performance>, 2014.
- [12] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4):547–553, 2009.
- [13] P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. Wine quality data set. <https://archive.ics.uci.edu/ml/datasets/Wine+Quality>, 2009.
- [14] P. Cortez and A. M. G. Silva. Using data mining to predict secondary school student performance. In *Future Business Technology Conference*, pages 5–12. EUROISIS, 2008.
- [15] I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In *CRYPTO*, pages 643–662. Springer, 2012.
- [16] S. de Hoogh, B. Schoenmakers, and M. Veeningen. Certificate validation in secure computation and its use in verifiable linear programming. In *AFRICACRYPT*, volume 9646 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2016.
- [17] D. Demmler, G. Dessouky, F. Koushanfar, A. Sadeghi, T. Schneider, and S. Zeitouni. Automated synthesis of optimized circuits for secure computation. In *ACM*

- Conference on Computer and Communications Security*, pages 1504–1517. ACM, 2015.
- [18] D. Demmler, T. Schneider, and M. Zohner. ABY - A framework for efficient mixed-protocol secure two-party computation. In *NDSS*. The Internet Society, 2015.
- [19] W. Du and M. J. Atallah. Privacy-preserving cooperative scientific computations. In *CSFW*, pages 273–294. IEEE Computer Society, 2001.
- [20] W. Du and M. J. Atallah. Protocols for secure remote database access with approximate matching. In *E-Commerce Security and Privacy*, volume 2 of *Advances in Information Security*, pages 87–111. Springer, 2001.
- [21] W. Du, Y. S. Han, and S. Chen. Privacy-preserving multivariate statistical analysis: Linear regression and classification. In *SDM*, pages 222–233. SIAM, 2004.
- [22] C. Dwork and K. Nissim. Privacy-preserving datamining on vertically partitioned databases. In *CRYPTO*, volume 3152 of *Lecture Notes in Computer Science*, pages 528–544. Springer, 2004.
- [23] C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [24] M. D. Ercegovac and T. Lang. *Digital arithmetic*. Elsevier, 2004.
- [25] H. Fanaee-T and J. Gama. Bike sharing dataset data set. <https://archive.ics.uci.edu/ml/datasets/Bike+Sharing+Dataset>, 2013.
- [26] H. Fanaee-T and J. Gama. Event labeling combining ensemble detectors and background knowledge. *Progress in AI*, 2(2-3):113–127, 2014.
- [27] J. Fonollosa and R. Huerta. Gas sensor array under dynamic gas mixtures data set. <https://archive.ics.uci.edu/ml/datasets/Gas+sensor+array+under+dynamic+gas+mixtures>, 2015.
- [28] J. Fonollosa, S. Sheik, R. Huerta, and S. Marco. Reservoir computing compensates slow response of chemosensor arrays exposed to fast varying gas concentrations in continuous monitoring. *Sensors and Actuators B: Chemical*, 215:618–629, 2015.
- [29] N. Gilboa. Two party RSA key generation. In *CRYPTO*, volume 1666 of *Lecture Notes in Computer Science*, pages 116–129. Springer, 1999.
- [30] O. Goldreich. *The Foundations of Cryptography - Volume 2, Basic Applications*. Cambridge University Press, 2004.
- [31] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *STOC*, pages 218–229. ACM, 1987.
- [32] S. D. Gordon, J. Katz, V. Kolesnikov, F. Krell, T. Malkin, M. Raykova, and Y. Vahlis. Secure two-party computation in sublinear (amortized) time. In *ACM Conference on Computer and Communications Security*, pages 513–524. ACM, 2012.
- [33] T. Graepel, K. E. Lauter, and M. Naehrig. ML confidential: Machine learning on encrypted data. In *ICISC*, pages 1–21. Springer, 2012.
- [34] F. Graf, H.-P. Kriegel, M. Schubert, S. Poelsterl, and A. Cavallaro. Relative location of ct slices on axial axis data set. <https://archive.ics.uci.edu/ml/datasets/Relative+location+of+CT+slices+on+axial+axis>, 2011.
- [35] R. Hall, S. E. Fienberg, and Y. Nardi. Secure multiple linear regression based on homomorphic encryption. *Journal of Official Statistics*, 27(4):669, 2011.
- [36] W. Henecka, S. Kögl, A. Sadeghi, T. Schneider, and I. Wehrenberg. TASTY: tool for automating secure two-party computations. In *ACM Conference on Computer and Communications Security*, pages 451–462. ACM, 2010.
- [37] Y. Huang, D. Evans, J. Katz, and L. Malka. Faster secure two-party computation using garbled circuits. In *USENIX Security Symposium*. USENIX Association, 2011.
- [38] Y. Huang, C. Shen, D. Evans, J. Katz, and A. Shelat. Efficient secure computation with garbled circuits. In *ICISS*, volume 7093 of *Lecture Notes in Computer Science*, pages 28–48. Springer, 2011.
- [39] A. Karatsuba and Y. Ofman. Multiplication of many-digital numbers by automatic computers. In *Proceedings of the USSR Academy of Sciences 145*, pages 293–294, 1962.
- [40] A. F. Karr, X. Lin, A. P. Sanil, and J. P. Reiter. Regression on distributed databases via secure multi-party computation. In *DG.O*, ACM International Conference Proceeding Series. Digital Government Research Center, 2004.
- [41] M. Keller, E. Orsini, and P. Scholl. MASCOT: faster malicious arithmetic secure computation with oblivious transfer. In *ACM Conference on Computer and Communications Security*, pages 830–842. ACM, 2016.
- [42] M. Keller and P. Scholl. Efficient, oblivious data structures for MPC. In *ASIACRYPT (2)*, pages 506–525. Springer, 2014.
- [43] D. E. Knuth. *The Art of Computer Programming, Volume 2 (3rd Ed.): Seminumerical Algorithms*. Addison-Wesley, 1997.
- [44] V. Kolesnikov and T. Schneider. Improved garbled circuit: Free XOR gates and applications. In *ICALP (2)*, pages 486–498. Springer, 2008.
- [45] P. Laud and M. Pettai. Secure multiparty sorting protocols with covert privacy. In *NordSec*, volume 10014 of *Lecture Notes in Computer Science*, pages 216–231, 2016.
- [46] M. Lichman. UCI machine learning repository. <http://archive.ics.uci.edu/ml>, 2013.
- [47] Y. Lindell. How to simulate it - A tutorial on the simulation proof technique. *IACR Cryptology ePrint Archive*, 2016:46, 2016.
- [48] Y. Lindell and B. Pinkas. Privacy preserving data mining. *J. Cryptology*, 15(3):177–206, 2002.
- [49] Y. Lindell and B. Pinkas. A proof of security of Yao’s protocol for two-party computation. *J. Cryptology*, 22(2):161–188, 2009.
- [50] Y. Lindell, B. Pinkas, N. P. Smart, and A. Yanai. Efficient constant round multi-party computation combining BMR and SPDZ. In *CRYPTO (2)*, pages 319–338. Springer, 2015.
- [51] C. Liu, X. S. Wang, K. Nayak, Y. Huang, and E. Shi. OblivM: A programming framework for secure computation. In *IEEE Symposium on Security and Privacy*, pages 359–376. IEEE Computer Society, 2015.
- [52] G. Meurant. *The Lanczos and conjugate gradient algorithms: from theory to finite precision computations*, volume 19. SIAM, 2006.
- [53] P. Mohassel and M. K. Franklin. Efficiency tradeoffs for malicious two-party computation. In *Public Key Cryptography*, volume 3958 of *Lecture Notes in Computer Science*, pages 458–473. Springer, 2006.

- [54] K. P. Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- [55] M. Naor and B. Pinkas. Efficient oblivious transfer protocols. In *SODA*, pages 448–457. ACM/SIAM, 2001.
- [56] M. Naor, B. Pinkas, and R. Sumner. Privacy preserving auctions and mechanism design. In *EC*, pages 129–139, 1999.
- [57] K. Nayak, X. S. Wang, S. Ioannidis, U. Weinsberg, N. Taft, and E. Shi. Graphsc: Parallel secure computation made easy. In *IEEE Symposium on Security and Privacy*, pages 377–394. IEEE Computer Society, 2015.
- [58] J. B. Nielsen, P. S. Nordholt, C. Orlandi, and S. S. Burra. A new approach to practical active-secure two-party computation. In *CRYPTO*, volume 7417 of *Lecture Notes in Computer Science*, pages 681–700. Springer, 2012.
- [59] V. Nikolaenko, U. Weinsberg, S. Ioannidis, M. Joye, D. Boneh, and N. Taft. Privacy-preserving ridge regression on hundreds of millions of records. In *IEEE Symposium on Security and Privacy*, pages 334–348. IEEE Computer Society, 2013.
- [60] J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [61] P. Pullonen and S. Siim. Combining secret sharing and garbled circuits for efficient private IEEE 754 floating-point computations. In *Financial Cryptography Workshops*, pages 172–183. Springer, 2015.
- [62] M. Rabin. How to Exchange Secrets by Oblivious Transfer. Technical Report TR-81, Harvard Aiken Computation Laboratory, 1981.
- [63] M. Redmond. Communities and crime data set. <https://archive.ics.uci.edu/ml/datasets/Communities+and+Crime>, 2009.
- [64] M. Redmond and A. Baveja. A data-driven software tool for enabling cooperative information sharing among police departments. *European Journal of Operational Research*, 141(3):660–678, 2002.
- [65] A. P. Sanil, A. F. Karr, X. Lin, and J. P. Reiter. Privacy preserving regression modelling via distributed computation. In *KDD*, pages 677–682. ACM, 2004.
- [66] E. M. Songhori, S. U. Hussain, A. Sadeghi, T. Schneider, and F. Koushanfar. Tinygarble: Highly compressed and scalable sequential garbled circuits. In *IEEE Symposium on Security and Privacy*, pages 411–428. IEEE Computer Society, 2015.
- [67] X. Wang, A. J. Malozemoff, and J. Katz. Faster secure two-party computation in the single-execution setting. In *EUROCRYPT (3)*, volume 10212 of *Lecture Notes in Computer Science*, pages 399–424, 2017.
- [68] M. H. Weik. A third survey of domestic electronic digital computing systems. Technical report, DTIC Document, 1961.
- [69] J. H. Wilkinson. *The algebraic eigenvalue problem*. Clarendon Press Oxford, 1988.
- [70] A. C. Yao. How to generate and exchange secrets (extended abstract). In *FOCS*, pages 162–167. IEEE Computer Society, 1986.
- [71] H. Yu, J. Vaidya, and X. Jiang. Privacy-preserving SVM classification on vertically partitioned data. In *PAKDD*, pages 647–656. Springer, 2006.
- [72] S. Zahur and D. Evans. Obliv-C: A language for extensible data-oblivious computation. *IACR Cryptology ePrint Archive*, 2015:1153, 2015.

- [73] S. Zahur, M. Rosulek, and D. Evans. Two halves make a whole - reducing data transfer in garbled circuits using half gates. In *EUROCRYPT (2)*, pages 220–250. Springer, 2015.
- [74] Absentminded crypto kit. <https://bitbucket.org/jackdoerner/absentminded-crypto-kit>.
- [75] Secure distributed linear regression. <https://github.com/schoppmp/linreg-mpc>.
- [76] Auto MPG data set. <https://archive.ics.uci.edu/ml/datasets/Auto+MPG>, 1993.

A Proof from Section 4.2

Theorem 3. *If $q > 2nR^2/\delta^2$, then:*

$$|\langle u, v \rangle - \tilde{\phi}_{\delta^2}(\tilde{\varphi}_q(\langle \vec{x}, \vec{y} \rangle))| \leq 2nR\delta + n\delta^2 .$$

Proof. In the first place we observe that any number occurring in the computation of the inner product $\langle \phi_{\delta}(\vec{u}), \phi_{\delta}(\vec{v}) \rangle$ of the integer encodings of u and v is bounded by nR^2/δ^2 . Therefore, the condition on q ensures there are no overflows in the computation in \mathbb{Z}_q and we have $\langle \phi_{\delta}(u), \phi_{\delta}(v) \rangle = \tilde{\varphi}_q(\langle \vec{x}, \vec{y} \rangle)$. Now we use the formulas for the error of sum and product of integer encodings from previous section to show that

$$\begin{aligned} & \left| \langle \vec{u}, \vec{v} \rangle - \tilde{\phi}_{\delta^2}(\langle \phi_{\delta}(\vec{u}), \phi_{\delta}(\vec{v}) \rangle) \right| \\ &= \left| \sum_{i=1}^n u_i v_i - \tilde{\phi}_{\delta^2} \left(\sum_{i=1}^n \phi_{\delta}(u_i) \phi_{\delta}(v_i) \right) \right| \\ &\leq \sum_{i=1}^n |u_i v_i - \tilde{\phi}_{\delta^2}(\phi_{\delta}(u_i) \phi_{\delta}(v_i))| \\ &\leq n(2R\delta + \delta^2) , \end{aligned}$$

where the first inequality follows by the triangle inequality and linearity of the decoding map $\tilde{\phi}_{\delta^2}$. \square

B Further Experimental Results




Fig. 7. Circuit sizes for 64 bits (left), 32 bits (right) as a function of the input dimensionality d . One can see a clear correlation between the circuit size and the execution time shown in Figure 6 (left).

n	d	Number of parties (data providers)					
		2		3		5	
		$b = 64$	$b = 32$	$b = 64$	$b = 32$	$b = 64$	$b = 32$
10000	20	24s	13s	20s	11s	14s	8s
	50	2m20s	1m16s	1m55s	1m2s	1m22s	43s
	100	9m6s	4m58s	7m18s	3m55s	5m13s	2m47s
20000	20	45s	24s	38s	20s	29s	15s
	50	4m26s	2m18s	3m39s	1m56s	2m39s	1m21s
	100	17m25s	9m58s	14m9s	7m24s	10m10s	5m14s
50000	20	1m50s	56s	1m32s	47s	1m7s	35s
	50	10m47s	5m27s	8m58s	4m37s	6m31s	3m12s
	100	42m12s	21m25s	34m39s	17m41s	24m58s	12m32s
100000	20	3m40s	1m50s	3m1s	1m33s	2m18s	1m7s
	50	21m40s	10m47s	17m25s	9m6s	13m5s	6m40s
	100	1h25m14s	42m18s	1h7m30s	35m3s	50m16s	24m12s
200000	20	7m25s	3m40s	5m57s	3m4s	4m46s	2m18s
	50	43m47s	21m43s	34m14s	18m26s	26m47s	12m40s
	100	2h52m8s	1h25m6s	2h11m35s	1h10m39s	1h43m12s	49m23s
500000	20	18m18s	9m10s	14m29s	7m27s	12m10s	5m52s
	50	1h47m59s	54m13s	1h23m49s	45m0s	1h8m8s	32m14s
	100	7h3m56s	3h32m37s	5h20m52s	2h53m54s	4h17m8s	2h4m53s

Table 4. Computation time of the aggregation phase using the OT-based inner product protocol, for 2, 3, and 5 data providers and different values of n (number of records) and d (number of features). Note that unlike the TI-based protocol in Table 5, this algorithm scales well with the number of parties.

n	d	Number of parties (data providers)								
		2			3			5		
		CSP	DP	Total	CSP	DP	Total	CSP	DP	Total
100000	20	1s	2s	3s	2s	2s	3s	2s	1s	4s
	50	7s	9s	14s	9s	8s	17s	11s	6s	20s
	100	27s	31s	50s	36s	28s	1m4s	43s	21s	1m15s
	200	1m46s	1m57s	3m12s	2m22s	1m44s	4m10s	2m50s	1m17s	4m54s
	500	10m52s	11m53s	19m14s	14m40s	10m29s	25m16s	17m29s	7m37s	30m2s
200000	20	2s	4s	6s	3s	4s	7s	4s	3s	8s
	50	14s	19s	30s	18s	17s	37s	22s	13s	42s
	100	55s	1m6s	1m47s	1m11s	59s	2m16s	1m26s	44s	2m39s
	200	3m37s	4m13s	7m0s	4m44s	3m43s	9m0s	5m38s	2m43s	10m32s
	500	21m46s	25m33s	41m47s	29m11s	22m23s	54m21s	34m56s	16m6s	1h4m0s
500000	20	6s	10s	15s	8s	9s	18s	10s	8s	21s
	50	36s	48s	1m17s	47s	42s	1m35s	57s	32s	1m51s
	100	2m17s	2m51s	4m47s	3m4s	2m29s	6m1s	3m42s	1m51s	6m58s
	200	9m0s	10m55s	18m52s	12m15s	9m33s	24m8s	14m47s	6m52s	27m43s
	500	56m22s	1h6m41s	1h53m52s	1h14m37s	57m31s	2h25m55s	1h30m41s	40m56s	2h50m53s
1000000	20	12s	20s	31s	16s	18s	37s	20s	15s	44s
	50	1m12s	1m34s	2m40s	1m37s	1m24s	3m21s	1m58s	1m4s	3m53s
	100	4m45s	5m40s	10m1s	6m16s	4m58s	12m42s	7m31s	3m41s	14m48s
	200	18m49s	21m44s	39m16s	25m7s	18m48s	49m56s	30m28s	13m48s	59m22s

Table 5. Computation time of the aggregation phase using the TI-based inner product protocol with 64 bits, for 2, 3, and 5 data providers and different values of n (number of records) and d (number of features). For each number of data providers, computation time of the CSP (left) and the data providers (middle, averaged) is reported, as well as total running time (right).