
Privacy-Preserving XGBoost Inference

Xianrui Meng Joan Feigenbaum*
AWS Cryptography Algorithm
{xianru, jfbaum}@amazon.com

Abstract

Although machine learning (ML) is widely used for predictive tasks, there are important scenarios in which ML cannot be used or at least cannot achieve its full potential. A major barrier to adoption is the sensitive nature of predictive queries. Individual users may lack sufficiently rich datasets to train accurate models locally but also be unwilling to send sensitive queries to commercial services that vend such models. One central goal of *privacy-preserving machine learning* (PPML) is to enable users to submit *encrypted queries* to a remote ML service, receive encrypted results, and decrypt them locally. We aim at developing practical solutions for real-world privacy-preserving ML inference problems. In this paper, we propose a privacy-preserving XGBoost prediction algorithm, which we have implemented and evaluated empirically on AWS SageMaker. Experimental results indicate that our algorithm is efficient enough to be used in real ML production environments.

1 Introduction

Machine Learning (ML) plays an important role in daily life. Pervasive use of digital devices and networks produces massive amounts of data that are analyzed to reveal patterns and correlations that, in turn, are used to draw conclusions or make predictions. Diverse applications that make successful use of ML include market forecasting, service personalization, voice and facial recognition, autonomous driving, health diagnostics, education, and security analytics.

Clearly in tension with the utility of ML is the desire of individuals and organizations for data privacy. Both the input and the output of an ML prediction may be highly personal, confidential information and may be constrained by regulations. For example, students' confidential educational records are governed by FERPA – the Family Educational Rights and Privacy Act. Organizations that produce valuable models may wish to sell access to them on a pay-per-prediction basis and must protect them as one would any valuable form of digital intellectual property. The need to maintain the privacy of data subjects, to protect intellectual property, and to keep commercially valuable instances and predictions confidential motivate the study of *privacy-preserving machine learning* (PPML).

Gentry has shown that one can perform arbitrary computations on encrypted data using fully homomorphic-encryption (FHE) (16). In principle, FHE could fully resolve the tension between utility of ML and data-privacy requirements, but there are application scenarios in which it is prohibitively computationally expensive by orders of magnitude. In these scenarios, it is natural to seek specialized homomorphic-encryption (SHE) schemes that are more efficient. We provide one such scheme in this paper.

Extreme Gradient Boosting (XGBoost) (9) is an optimized, distributed, gradient-boosting ML framework designed to be highly efficient, flexible, and portable. It performs *parallel tree boosting* that solves many classification and regression problems quickly and accurately. For example, 17 of the 29 challenge-winning solutions published on Kaggle's blog in 2015 used XGBoost (33), Facebook uses

*Amazon Scholar and Grace Hopper Prof. of Computer Science at Yale University

it to predict click through on advertisements (19), and it is very popular on Amazon’s managed-cloud ML platform SageMaker.²

We present PPXGBoost, a privacy-preserving XGBoost-prediction algorithm, in Section 2. In Section 3, we explain the security definition and the privacy properties that our algorithm achieves. Experimental results are given in Section 4. Finally, we present open problems in Section 5.

2 Privacy-preserving XGBoost

2.1 Preliminaries

XGBoost Upon receiving a training dataset, the XGBoost training algorithm produces an ML model in the form of a set $\{T_i\}_{i=1}^M$ of *classification and regression trees* (CARTs). A CART is a generalization of a decision tree; while the latter produces a binary output, thus classifying each input query as a “yes” or “no” instance of the phenomenon under study, a CART assigns to each input query a (real) numerical score. Interpretation of scores is application-dependent.³ We use boldface lowercase letters, such as \mathbf{v} , to denote a vector of real numbers. If \mathbf{v} is an input query, and $\{y_i \leftarrow T_i(\mathbf{v})\}_{i=1}^M$ is the set of scores produced by the CARTs, then the final prediction (*i.e.*, the overall score assigned to \mathbf{v} by XGBoost) is typically $y = \sum_{i=1}^M y_i$. Depending on the application, a *softmax* function may be applied to the y_i to obtain the final prediction, but we restrict attention to summation in this paper. A full explanation of XGBoost training and inference can be found in the original paper of Chen and Guestrin (9).

Homomorphic encryption Homomorphic encryption is a form of encryption that can perform arbitrary computation on plaintext values while manipulating only ciphertexts. In this work, we use an additive SHE scheme. Specifically, let $\text{SHE} = (\text{Gen}, \text{Enc}, \text{Dec})$ be a public-key SHE that consists of three polynomial-time algorithms. Gen is a probabilistic algorithm that takes a security parameter k as input and returns a private and public key pair (pk, sk) . Enc is a probabilistic encryption algorithm that takes as input a public key pk and a message m and outputs a ciphertext. Dec is an algorithm that takes as input a private key sk and a ciphertext and returns the corresponding plaintext. In addition, SHE has an evaluation algorithm Eval that supports any number of additions over the plaintexts: $\text{Enc}(\text{pk}, (m_1 + \dots + m_n)) = \text{Eval}(+, \text{Enc}(\text{pk}, m_1), \dots, \text{Enc}(\text{pk}, m_n))$.

Order-preserving encryption For $A, B \subseteq \mathcal{N}$ with $|A| \leq |B|$, a function $f : A \rightarrow B$ is *order-preserving* if, for all $i, j \in A$, $f(i) > f(j)$ if and only if $i > j$. We say that a symmetric encryption scheme $(\text{Gen}, \text{Enc}, \text{Dec})$ with plaintext and ciphertext spaces \mathcal{D} and \mathcal{R} is an *order-preserving encryption* (OPE) scheme if $\text{Enc}(K, \cdot)$ is an order-preserving function from \mathcal{D} to \mathcal{R} , for all K output by $\text{Gen}(1^k)$. To make sense of the $>$ relation in this context, elements of \mathcal{D} and \mathcal{R} are encoded as binary strings, which are then interpreted as numbers.

Throughout this paper, “polynomial” means “polynomial in the security parameter k .” Formal definitions of these cryptographic concepts, including the Pseudorandom Function (PRF) family and semantic security can be found in (23).

2.2 PPXGBoost inference algorithm

The PPXGBoost algorithm is given in Figure 1. On the client side, there is an app with which a user encrypts queries and decrypts results. On the server side, there is a module called Proxy that runs in a trusted environment and is responsible for set up (*i.e.*, creating, for each authorized user, an encrypted model and a set of cryptographic keys) and an ML module that evaluates the encrypted queries.

The inputs to the Setup phase are an XGBoost model Ω , a model hyper-parameter α , and a security parameter k . During this phase, Proxy generates, for each authorized user, the keys K_1 and (pk, sk) . K_1 is the user’s key for the (symmetric-key) OPE scheme. (pk, sk) is the user’s key pair for the (public-key) SHE scheme. Proxy then encrypts the node values in each CART T_i in Ω to create an encrypted CART T'_i in this user’s encrypted model EncML . For each internal node in T_i with value \mathbf{x} , the value of the corresponding node in T'_i is $\text{OPE.Enc}(K_1, \mathbf{x})$. (Vectors of values are encrypted

²<https://aws.amazon.com/sagemaker/>

³This is one of the roles of the hyper-parameter α referred to in Subsection 2.2

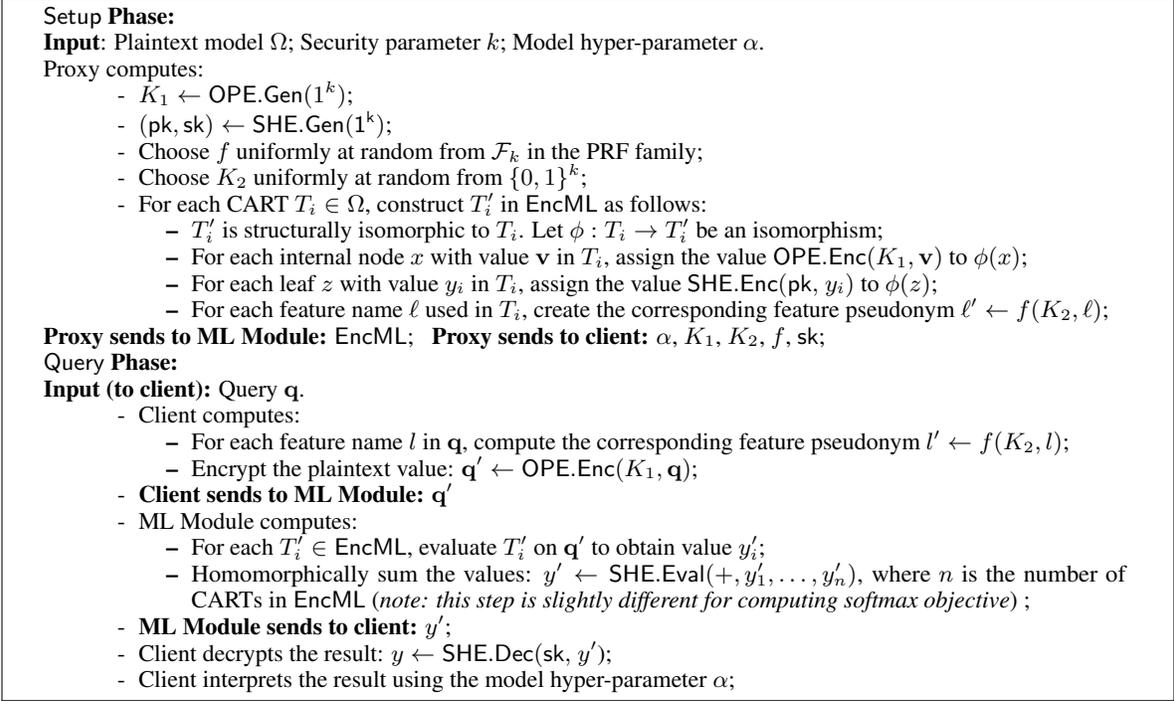


Figure 1: PPXGBoost: A privacy-preserving XGBoost inference algorithm

and decrypted component-wise.) For each leaf in T_i with value y , the value of the corresponding leaf in T'_i is $\text{SHE.Enc}(\text{pk}, y)$. Finally, the proxy sends K_1, K_2, f and sk to the user’s client and sends EncML to the ML module.

In the Query phase, the client first encrypts its plaintext query \mathbf{q} with OPE, *i.e.*, it computes $\mathbf{q}' \leftarrow \text{OPE.Enc}(K_1, \mathbf{q})$. It sends \mathbf{q}' to the ML module, which evaluates each T'_i in EncML on input \mathbf{q}' to obtain a value y_i . The module computes $y' \leftarrow \text{SHE.Eval}(+, y'_1, \dots, y'_n)$, where n is the number of CARTs, and sends it to the client, which decrypts to obtain the final result $y \leftarrow \text{SHE.Dec}(\text{sk}, y')$.

The correctness of this scheme follows directly from the properties of OPE and SHE. Because $a > b$ if and only if $\text{OPE.Enc}(K_1, a) > \text{OPE.Enc}(K_1, b)$, for all a, b , and K_1 , and the same K_1 is used to encrypt both queries and internal-node values in Ω , an encrypted query will travel precisely the same path through each encrypted CART T'_i that the corresponding plaintext query would have traveled through the corresponding plaintext CART T_i . Because the leaf values y_i in Ω have been encrypted using the additively homomorphic encryption operation $\text{SHE.Enc}(\text{pk}, y_i)$, and sk is the decryption key that corresponds to pk , the plaintext y corresponding to the ciphertext sum y' is the sum of the individual plaintext values y_i in leaves of T_i .

The proxy also chooses, for each authorized user, a function f uniformly at random from \mathcal{F}_k in the PRF family and a key K_2 uniformly at random from $\{0, 1\}^k$ for use with f . This function is used to generate pseudorandom “feature names” for vectors of queries and node values. We defer discussion of this aspect of the algorithm until the full paper.

Note that the plaintext (Ω, α, k) can be used by a very large user population, but a unique, personalized encrypted model must be created for each individual user.

3 Privacy properties

Ideally, we would like a privacy-preserving inference algorithm to hide all information about the model, the queries, and the results from all adaptive probabilistic polynomial-time (PPT) adversaries.

For the Setup phase, this means that the adversary should be able to choose a sequence M_1, M_2, \dots, M_n of plaintext models, submit them to an oracle, and receive the corresponding sequence M'_1, M'_2, \dots, M'_n of encrypted models; it may choose the sequence adaptively in the sense that its choice of M_i may depend upon the oracle’s answers M'_1, \dots, M'_{i-1} . After this adaptive, chosen-plaintext,

oracle-query phase, the adversary is presented with an encrypted model that it has not seen before, and it cannot infer anything about the corresponding plaintext model.

For the Query phase, this means that, for a fixed encrypted model M' , the adversary should be able to choose a sequence q_1, q_2, \dots, q_n of plaintext queries, submit them to an oracle, and receive the corresponding sequence $\sigma = (q'_1, r'_1), (q'_2, r'_2), \dots, (q'_n, r'_n)$ of encrypted queries and encrypted results; once again, it may choose q_i based on $\sigma_{i-1} = (q'_1, r'_1), \dots, (q'_{i-1}, r'_{i-1})$. After this adaptive query phase, it cannot infer anything about the encrypted model; furthermore, when subsequently presented with additional pairs $(q'_{n+1}, r'_{n+1}), \dots, (q'_{n+j}, r'_{n+j})$, it cannot infer anything about the corresponding plaintext queries or answers.

Known algorithms that achieve these ideal privacy properties are not efficient enough for practical use. As initiated in the work of Curtmola *et al.* (10) and Chase and Kamara (8), one can instead define acceptable *leakage functions* and devise efficient algorithms that *provably* leak only the values of these functions. In PPXGBoost, this information may be leaked to the ML module and any party that observes the inner workings of the ML module, the communication between Proxy and the ML module, or the communication between the client and the ML module.

Because of space limitations we give the main ideas of our formal security definitions, PPXGBoost’s privacy properties, and our security proof in Appendix B.

4 AWS SageMaker experiments

DATASET	TIME		MODEL SIZE	
	XGBoost	PPXGBoost	XGBoost	PPXGBoost
AMAZON SYNTHETIC DATA	1ms	0.43s	506KB	4.2MB
TITANIC	< 1ms	0.32s	3KB	12KB
US CENSUS	1ms	0.49s	210KB	2.5MB

Table 1: PPXGBoost Performance

For our experiments on PPXGBoost, we implemented the cryptographic protocols in python3. We instantiated the PRF using HMAC, and we used Paillier encryption (31) for our additive SHE and Boldyreva *et al.*’s scheme (3) for our OPE.

An overview of our system architecture can be found in Appendix A. All of our experiments were run on AWS. The Setup phase is deployed in AWS Virtual Private Cloud ⁴ environment. The inference procedure is run on SageMaker using an ml.t2.large instance and the Amazon Elastic Container Service ⁵. Our experimental results are summarized in Table 1. We ran PPXGBoost on three different models derived from three different datasets. One data is synthetically generated based on Amazon’s dataset. The other two datasets are public datasets. On average, PPXGBoost inference is approximately 10^3 times slower than the plaintext version of XGBoost. The size of encrypted models is between four and nine times larger than that of the plaintext models. The inference time includes the network traffic time. This performance is sufficient for many inference tasks currently done on smart phones that must query a remote server.

5 Open problems

Our initial version of PPXGBoost is still quite limited. Currently, we can support binary classifications and multiclass classification using the softmax objective. Future work includes support for more learning parameters in the privacy-preserving version. Moreover, in our algorithm, we leverage the order-preserving encryption scheme to support comparisons. Comparison on semantically encrypted data is computationally expensive, but we plan to investigate the use of secure multiparty computation for this purpose. In particular, we will explore the use of two non-colluding servers that execute secure comparison for each internal node in an encrypted CART.

⁴<https://aws.amazon.com/vpc/>

⁵<https://aws.amazon.com/ecs/>

References

- [1] Yoshinori Aono, Takuya Hayashi, Le Trieu Phong, and Lihua Wang. Scalable and secure logistic regression via homomorphic encryption. In *6th ACM Conference on Data and Application Security and Privacy*, pages 142–144, 2016.
- [2] Mauro Barni, Pierluigi Failla, Riccardo Lazzeretti, Ahmad-Reza Sadeghi, and Thomas Schneider. Privacy-preserving ECG classification with branching programs and neural networks. *IEEE Trans. Information Forensics and Security*, 6(2):452–468, 2011.
- [3] Alexandra Boldyreva, Nathan Chenette, and Adam O’Neill. Order-preserving encryption revisited: improved security analysis and alternative solutions. In *31st CRYPTO*, volume 6841 of *Lecture Notes in Computer Science*, pages 578–595. Springer, 2011.
- [4] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *24th ACM Conference on Computer and Communications Security*, pages 1175–1191, 2017.
- [5] Raphael Bost, Raluca Ada Popa, Stephen Tu, and Shafi Goldwasser. Machine learning classification over encrypted data. In *22nd Annual Network and Distributed System Security Symposium*. The Internet Society, 2015.
- [6] Hervé Chabanne, Amaury de Wargny, Jonathan Milgram, Constance Morel, and Emmanuel Prouff. Privacy-preserving classification on deep neural network. Cryptology ePrint Archive, Report 2017/035, 2017. <https://eprint.iacr.org/2017/035>.
- [7] Nishanth Chandran, Divya Gupta, Aseem Rastogi, Rahul Sharma, and Shardul Tripathi. Ezpc: Programmable, efficient, and scalable secure two-party computation. Cryptology ePrint Archive, Report 2017/1109, 2017. <https://eprint.iacr.org/2017/1109>.
- [8] Melissa Chase and Seny Kamara. Structured encryption and controlled disclosure. In *16th ASIACRYPT*, volume 6477 of *Lecture Notes in Computer Science*, pages 577–594. Springer, 2010.
- [9] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.
- [10] Reza Curtmola, Juan Garay, Seny Kamara, and Rafail Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *13th ACM Conference on Computer and Communications Security*, pages 79–88, 2006.
- [11] Wenliang Du and Mikhail J. Atallah. Privacy-preserving cooperative scientific computations. In *14th IEEE Computer Security Foundations Workshop*, pages 273–294, 2001.
- [12] Wenliang Du, Yunghsiang S. Han, and Shigang Chen. Privacy-preserving multivariate statistical analysis: Linear regression and classification. In *4th SIAM International Conference on Data Mining*, pages 222–233, 2004.
- [13] Stephen E. Fienberg, William J. Fulp, Aleksandra B. Slavkovic, and Tracey A. Wrobel. “Secure” log-linear and logistic regression analysis of distributed databases. In *Privacy in Statistical Databases*, volume 4302 of *Lecture Notes in Computer Science*, pages 277–290. Springer, 2006.
- [14] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *22nd ACM Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- [15] Adrià Gascón, Phillipp Schoppmann, Borja Balle, Mariana Raykova, Jack Doerner, Samee Zahur, and David Evans. Privacy-preserving distributed linear regression on high-dimensional data. *Proceedings on Privacy-Enhancing Technologies*, 2017(4):345–364, 2017.
- [16] Craig Gentry. Fully homomorphic encryption using ideal lattices. In *41st ACM Symposium on Theory of Computing*, pages 169–178, 2009.

- [17] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 201–210, 2016.
- [18] Ran Gilad-Bachrach, Kim Laine, Kristin E. Lauter, Peter Rindal, and Mike Rosulek. Secure data exchange: A marketplace in the cloud. In *10th ACM Conference on Cloud Computing Security Workshop*, pages 117–128, 2019.
- [19] Xinran He, Junfeng Pan, Ou Jin, Tianbing Xu, Bo Liu, Tao Xu, Yanxin Shi, Antoine Atallah, Ralf Herbrich, Stuart Bowers, and Joaquin Quiñero Candela. Practical lessons from predicting clicks on ads at Facebook. In *8th ACM International Workshop on Data Mining for Online Advertising*, pages 1–9, 2014.
- [20] Ehsan Hesamifard, Hassan Takabi, and Mehdi Ghasemi. Cryptodl: Deep neural networks over encrypted data. *Computing Research Repository*, abs/1711.05189, 2017.
- [21] Briland Hitaj, Giuseppe Ateniese, and Fernando Pérez-Cruz. Deep models under the GAN: information leakage from collaborative deep learning. In *24th ACM Conference on Computer and Communications Security*, pages 603–618, 2017.
- [22] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. Gazelle: A low latency framework for secure neural network inference. *Computing Research Repository*, abs/1801.05507, 2018.
- [23] Johnathan Katz and Yehuda Lindell. *Introduction to Modern Cryptography*. Chapman & Hall/CRC, 2008.
- [24] Yi Li, Yitao Duan, and Wei Xu. Privpy: Enabling scalable and general privacy-preserving computation. *Computing Research Repository*, abs/1801.10117, 2018.
- [25] Yehuda Lindell and Benny Pinkas. Privacy-preserving data mining. In *20th CRYPTO*, volume 1880 of *Lecture Notes in Computer Science*, pages 36–54. Springer, 2000.
- [26] Jian Liu, Mika Juuti, Yao Lu, and N. Asokan. Oblivious neural network predictions via minionn transformations. In *24th ACM Conference on Computer and Communications Security*, pages 619–631, 2017.
- [27] Yang Liu, Zhuo Ma, Ximeng Liu, Siqi Ma, Surya Nepal, and Robert H. Deng. Boosting privately: Privacy-preserving federated extreme boosting for mobile crowdsensing. *Computing Research Repository*, abs/1907.10218, 2019.
- [28] Payman Mohassel and Yupeng Zhang. SecureML: A system for scalable privacy-preserving machine learning. In *38th IEEE Symposium on Security and Privacy*, pages 19–38, 2017.
- [29] Valeria Nikolaenko, Udi Weinsberg, Stratis Ioannidis, Marc Joye, Dan Boneh, and Nina Taft. Privacy-preserving ridge regression on hundreds of millions of records. In *34th IEEE Symposium on Security and Privacy*, pages 334–348, 2013.
- [30] Claudio Orlandi, Alessandro Piva, and Mauro Barni. Oblivious neural network computing via homomorphic encryption. *EURASIP Journal on Information Security*, 2007:Article ID 37343.
- [31] Pascal Paillier. Public-key cryptosystems based on composite degree residuosity classes. In *7th EUROCRYPT*, volume 1592 of *Lecture Notes in Computer Science*, pages 223–238. Springer, 1999.
- [32] Manoj Prabhakaran and Amit Sahai, editors. *Secure Multi-Party Computation*, volume 10 of *Cryptology and Information Security Series*. IOS Press, 2013.
- [33] Kaggle Competition Ranking. <https://www.kaggle.com/c/otto-group-product-classification-challenge/discussion/14335>.
- [34] M. Sadegh Riazi, Christian Weinert, Oleksandr Tkachenko, Ebrahim M. Songhori, Thomas Schneider, and Farinaz Koushanfar. Chameleon: A hybrid secure computation framework for machine learning applications. *Computing Research Repository*, abs/1801.03239, 2018.

- [35] Bitza Darvish Rouhani, M. Sadegh Riazi, and Farinaz Koushanfar. Deepsecure: Scalable provably-secure deep learning. *Computing Research Repository*, abs/1705.08963, 2017.
- [36] Ahmad-Reza Sadeghi and Thomas Schneider. Generalized universal circuits for secure evaluation of private functions with application to data classification. In *11th International Conference on Information Security and Cryptology*, pages 336–353, 2008.
- [37] Ashish P. Sanil, Alan F. Karr, Xiaodong Lin, and Jerome P. Reiter. Privacy preserving regression modelling via distributed computation. In *10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 677–682, 2004.
- [38] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *38th IEEE Symposium on Security and Privacy*, pages 3–18, 2017.
- [39] Jaideep Vaidya, Hwanjo Yu, and Xiaoqian Jiang. Privacy-preserving SVM classification. *Knowledge and Information Systems*, 14(2):161–178, 2008.

A System Architecture

As mentioned in the paper, we deploy PPXGBoost using AWS infrastructure. We set up an Amazon VPC environment for deploying the inference prototype. Amazon VPC environment allows the model provider to have a logically isolated section of the AWS Cloud, therefore; the PPXGBoost provider can have complete control over the virtual networking environment. The Proxy service is deployed in a trusted environment, similarly to Amazon’s Key Management Services (KMS). The ML module is run on the Amazon SageMaker platform, a fully managed machine learning service. The security of SageMaker its own relies on the traditional AWS’s security model ⁶, including AWS Identity and Access Management (IAM), Amazon Macie, etc.

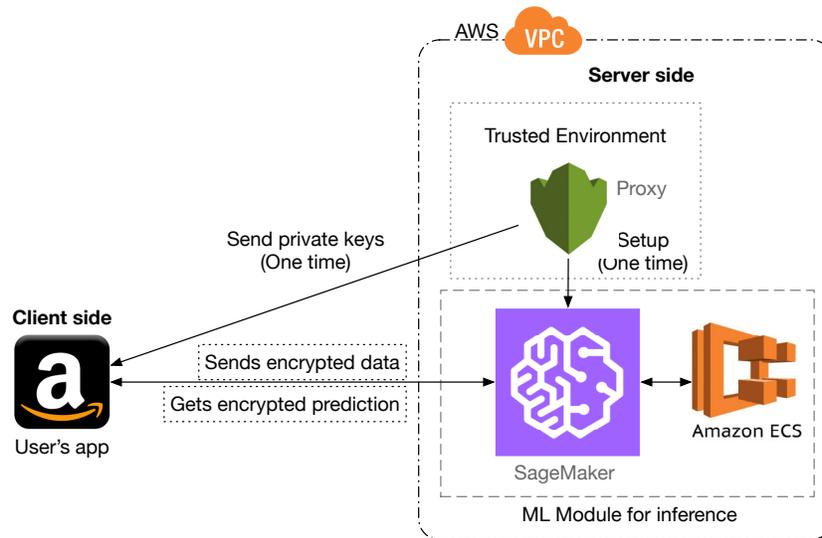


Figure 2: System architecture

Amazon SageMaker makes extensive use of Docker containers for build and runtime tasks. After the Proxy produces an encrypted model EncML, we store EncML to an S3 bucket (with proper permission configuration). We package our inference algorithm using Amazon’s container service, Amazon ECS. When deploying the ML inference module, we upload the inference package in Amazon’s Elastic Container Repository and specify the encrypted model location in S3. We create an endpoint for this SageMaker instance to handle the encrypted queries. After the client receives a private key, the client can send an encrypted query by querying a SageMaker endpoint. The ML module computes the encrypted query and returns an encrypted result to the client.

⁶see <https://docs.aws.amazon.com/sagemaker/latest/dg/data-protection.html>

B Privacy Definitions and Proof

At a high level, the security guarantee we require from privacy-preserving inference scheme is that: (1) given an encrypted ML model, no adversary can learn any information about the model; and (2) given the view of a polynomial number of Query executions for an adaptively generated sequence of queries $\mathbf{q} = (q_1, \dots, q_m)$, no adversary can learn any partial information about either the model or \mathbf{q} .

Such a security notion can be difficult to achieve efficiently, so often one allows for some form of leakage. Following (10; 8), this is usually formalized by parameterizing the security definition with leakage functions for each operation of the scheme which in this case include the Setup algorithm and Query protocol.

B.1 Security definition

In our description of what it means for $\text{PPXGBoost} = (\text{Setup}, \text{Query})$ to be secure, \mathcal{A} is a semi-honest adversary, \mathcal{S} is a simulator, $\mathcal{L}_{\text{Setup}}$ and $\mathcal{L}_{\text{Query}}$ are the leakage functions, and σ and σ_{i-1} are as defined above. The terms *Ideal* and *Real* are used as they are in the literature on searchable encryption (10).

Let Ω and α be an XGBoost model and hyper-parameter chosen by \mathcal{A} . Let v be a polynomially bounded function of k . We consider the following two randomized experiments.

Real($\Omega, 1^k, \alpha$):

- Run the Setup protocol: $(K, \text{EncML}) \leftarrow \text{Setup}(\Omega, 1^k, \alpha)$
- Execute the Query protocol $m = v(k)$ times. In these executions, Client (adaptively) chooses queries as well as performing its role in Figure 1; it chooses q_1 uniformly at random.
- Output the sequence $\sigma = (q'_1, r'_1), (q'_2, r'_2), \dots, (q'_m, r'_m)$ of encrypted queries and results.

Ideal $_{\mathcal{A}, \mathcal{S}}$ ($\mathcal{L}_{\text{Setup}}(\Omega), 1^k, \alpha$):

- Given $\mathcal{L}_{\text{Setup}}(\Omega), 1^k$, and α , \mathcal{S} generates an encrypted model EncML and sends it to \mathcal{A} .
- \mathcal{A} and \mathcal{S} conduct $m = v(k)$ executions of the Query protocol, in which \mathcal{S} plays the role of Client by (adaptively) constructing a sequence (q_1, \dots, q_m) of queries, and \mathcal{A} plays the role of the ML module. \mathcal{S} generates q_1 uniformly at random and uses $\mathcal{L}_{\text{Query}}(\Omega, \sigma_{i-1})$ to generate $q_i, 2 \leq i \leq m$.
- Output the sequence $\sigma = (q'_1, r'_1), (q'_2, r'_2), \dots, (q'_m, r'_m)$ of encrypted queries and results.

The gist of our security definition is that a PPT observer cannot distinguish between outputs of the **Real** experiment, which runs the protocol in Figure 1, and the **Ideal** experiment, in which a simulator that knows the values of the leakage functions plays the role of the client and the adversary plays the role of the ML module. A *distinguisher* \mathcal{D} is an algorithm that plays a refereed game with the adversary. In each round of the game, the referee obtains Ω and α from \mathcal{A} , runs either **Real**($\Omega, 1^k, \alpha$) or **Ideal** $_{\mathcal{A}, \mathcal{S}}$ ($\mathcal{L}_{\text{Setup}}(\Omega), 1^k, \alpha$), and shows the output of whichever experiment is run to the distinguisher.

Let w be a polynomially bounded function of k . The entire *distinguishing game* proceeds as follows. Fix a security parameter k . For $w(k)$ rounds, \mathcal{D} may ask the referee to run either **Real**($\Omega, 1^k, \alpha$) or **Ideal** $_{\mathcal{A}, \mathcal{S}}$ ($\mathcal{L}_{\text{Setup}}(\Omega), 1^k, \alpha$) and show him the output. The referee then chooses $b \in \{0, 1\}$ uniformly at random; if $b = 0$, it runs **Real**($\Omega, 1^k, \alpha$), and, if $b = 1$, it runs **Ideal** $_{\mathcal{A}, \mathcal{S}}$ ($\mathcal{L}_{\text{Setup}}(\Omega), 1^k, \alpha$). It shows the output to \mathcal{D} , and \mathcal{D} outputs its best guess b' of the value of b . The *distinguisher's advantage* in this game, which is a function of the security parameter k , is $|\Pr [b = b'] - \frac{1}{2}|$.

Definition B.1. We say that PPXGBoost is *adaptively* ($\mathcal{L}_{\text{Setup}}, \mathcal{L}_{\text{Query}}$)-*semantically secure* if, for every PPT adversary \mathcal{A} , there exists a PPT simulator \mathcal{S} for which every PPT distinguisher's advantage is negligible in k .

B.2 Leakage profile

We now describe the leakage functions that specify, in the sense of (10; 8), the information that PPXGBoost is willing to leak for the sake of efficiency.

Setup leakage. Recall that the Setup phase of PPXGBoost takes as one of its inputs a plaintext model Ω and gives as one of the outputs an encrypted model. The plaintext model consists of a set of CARTs. Setup leakage in PPXGBoost is a function $\mathcal{L}_{\text{Setup}}(\Omega)$ of the plaintext model; it consists of the number of CARTs in Ω , the depth of each CART, and, for each internal node w in each CART, which of w 's two children has the smaller value. Note that the numerical values of the nodes are *not* leaked; this is true of both internal nodes and leaves. In the high-level descriptions of Setup given in Section 2 and Figure 1, the entire structure of each CART is leaked, but, in practice, it is straightforward to pad each CART out to a complete binary tree of the appropriate depth without changing the results of the computation.

Query leakage. During the Query phase of PPXGBoost, the client and ML module exchange a sequence $\sigma = (q'_1, r'_1), (q'_2, r'_2), \dots, (q'_n, r'_n)$ of encrypted queries and encrypted results. Query leakage in PPXGBoost is a function $\mathcal{L}_{\text{Query}}(\Omega, \sigma)$ of the plaintext model and this sequence. It consists of a *query pattern* and the *set of paths* that are traversed during the execution of the encrypted queries. For every encrypted query q' in σ , this phase of PPXGBoost leaks the number of times it appears in σ and where it appears; that is, for every q' , the query phase reveals the set of i , $1 \leq i \leq n$, such that $q'_i = q'$. In addition, for each q'_i and each encrypted CART, the path from the root to a leaf in that CART that is traversed during the evaluation of q'_i is leaked to the ML module and to any party that can observe the inner workings of this module while queries are executed. Note that the query pattern and set of paths is well defined for each prefix σ_i of σ . Crucially, the decryptions of the queries and results are *not* leaked.

B.3 Main idea and interpretation of the security proof

To prove that the *only* information leaked by PPXGBoost is $\mathcal{L}_{\text{Setup}}$ and $\mathcal{L}_{\text{Query}}$, we present a PPT algorithm that is given 1^k , α , $\mathcal{L}_{\text{Setup}}(\Omega)$, and $\mathcal{L}_{\text{Query}}(\Omega, \sigma)$ as input and *simulates* the behavior of PPXGBoost's Setup and Query phases. We provide the main idea of the security proof here and defer the full proof to an expanded version of this paper.

Given $\mathcal{L}_{\text{Setup}}(\Omega)$, a simulator \mathcal{S} can construct a set $\{T_i\}$ of CARTs with the required ordering of internal nodes by sampling random values from the co-domain of OPE. For the leaves, \mathcal{S} assigns values chosen at random from the co-domain of SHE. Given $\mathcal{L}_{\text{Query}}(\Omega, \sigma)$, \mathcal{S} follows, for each encrypted query, the appropriate paths in each T_i that it constructed and returns the leaf value. The security properties of the OPE and SHE schemes ensure that the final predications are not revealed.

C Related work

Practical attacks on supervised learning systems that result in leakage of sensitive information about training datasets, models, or hyper-parameters can be found in, *e.g.* (14; 21; 38). Among proposals to mitigate those attacks, the majority focus on classification models, including decision trees (25), SVM classification (39), linear regression (11; 12; 37), logistic regression (13), and neural networks (30; 2; 36). Recently, a fast-growing number of works (*e.g.*, (5; 29; 15; 18; 1; 6; 34; 35; 26; 7; 20; 17; 28; 24; 22)) have achieved strong security guarantees in this setting by providing concrete security definitions and provably secure protocols that use multiple cryptographic-computation techniques (32). Another research thread has focused on privacy-preserving federated learning (*see, e.g.*, (4)), in which multiple mobile users update a global model by sharing aggregated updates to model parameters using a privacy-preserving, client-server protocol. Recently, Liu *et al.* (27) proposed a privacy-preserving boosting method for training XGBoost models in the federated-learning setting.