

Blind Faith: Privacy-Preserving Machine Learning using Function Approximation

Tanveer Khan

department of computing sciences
Tampere university
Tampere, Finland
tanveer.khan@tuni.fi

Alexandros Bakas

department of computing sciences
Tampere university
Tampere, Finland
alexandros.bakas@tuni.fi

Antonios Michalas

department of computing sciences
Tampere university
Tampere, Finland
antonios.michalas@tuni.fi

Abstract—Over the past few years, a tremendous growth of machine learning was brought about by a significant increase in adoption of cloud-based services. As a result, various solutions have been proposed in which the machine learning models run on a remote cloud provider. However, when such a model is deployed on an untrusted cloud, it is of vital importance that the users’ privacy is preserved. To this end, we propose Blind Faith – a machine learning model in which the training phase occurs in plaintext data, but the classification of the users’ inputs is performed on homomorphically encrypted ciphertexts. To make our construction compatible with homomorphic encryption, we approximate the activation functions using Chebyshev polynomials. This allowed us to build a privacy-preserving machine learning model that can classify encrypted images. Blind Faith preserves users’ privacy since it can perform high accuracy predictions by performing computations directly on encrypted data.

Index Terms—Neural Networks, Homomorphic Encryption, Privacy, Polynomial Approximation, Activation Functions

I. INTRODUCTION

Machine Learning (ML) models have attracted global adulation and are used in a plethora of applications such as medical diagnosis, pattern recognition, and credit risk assessment. Recently, deep learning – a sub-field of ML has gained extra attention from researchers due to its solid performance in many tasks such as speech recognition, spam detection, image classification, traffic analysis, face recognition, financial detection and genomics prediction [1]–[5]. ML models normally consist of a training and a testing phase. In the training phase, the model is trained on a dataset while in the testing phase, the model predicts the output using a previously unseen input. Due to the increased demand for these services, Cloud Service Providers (CSP) also offer Machine Learning as a Service (MLaaS). In MLaaS, the training and deploying of ML models are performed using the CSP infrastructure. Once the models are trained and deployed on the CSP, the users can use these models for online prediction services. MLaaS have clear benefits and such services are currently offered by providers including Google Prediction API [6], Microsoft Azure ML [7],

and Ersatz Lab [8]. The problem with MLaaS is that either the training or testing, or both phases, need to be outsourced to the CSP. The outsourced data can be of a sensitive nature, as for example in the areas of healthcare and finance [9], which, in its turn, raises privacy concerns. As an example, the data sent to the prediction models may be misused or stolen. To preserve the privacy of user data, different methods have been proposed [3], [10], [11].

The purpose of this work is to illustrate the use of Neural Networks (NN) over encrypted data. We use Homomorphic Encryption (HE) to perform computations on encrypted data. HE can perform arithmetic operations (addition and multiplication) over encrypted data without decrypting it. Therefore, any function that uses these arithmetic operations can be homomorphically evaluated. Furthermore, we consider the Convolutional Neural Network (CNN) to be homomorphically evaluated on the encrypted data. All the operations except activation functions in CNN are simple addition and multiplication and can be homomorphically evaluated. The activation functions such as rectified-linear unit (ReLU), Sigmoid, etc. are non-linear and cannot be evaluated homomorphically.

For CNN to be homomorphically evaluated on encrypted data, designing efficient HE-friendly support for activation function has been an active topic of research. Different methods are proposed to support the non-linear activation functions such as power functions [3], look-up table [12] and polynomial approximations [13]–[15]. In this work, we use low degree Chebyshev polynomials to approximate a non-linear activation function.

Chebyshev polynomials allow us to efficiently compute any continuous function in a given interval, using only low degree polynomials. This is an important feature as our computations become significantly more efficient and lower the overall computational complexity.

A. Our Contribution

The contributions of this paper are manifold. First, we show how to approximate activation functions using Chebyshev polynomials. As a next step, we used the approximated functions to design an NN model. Moreover, we combined our

Funded by the ASCLEPIOS: Advanced Secure Cloud Encrypted Platform for Internationally Orchestrated Solutions in Healthcare Project No. 826093 EU research project.

design with HE in order to deploy Blind Faith – a privacy-preserving ML (PPML) model that can classify encrypted images with high accuracy. Our model has a wide application domain since it can be applied to fields where users’ data is considered sacrosanct. For example, consider a cloud-based eHealth service where patients can upload an encrypted medical image and the cloud can approximately tell if the user suffers from a certain disease or not without getting any information about the actual user, the data sent to the cloud or any personal medical records. To illustrate the effectiveness of our model, we conducted extensive experiments. Our contributions can be summarized as follows:

- We approximate the ReLU and Sigmoid activation functions using Chebyshev polynomials and compare the differences in terms of efficiency and accuracy.
- We design a PPML model in which the CNN is trained on plaintext data while the classification process is performed on homomorphically encrypted data.
- We compare our design with other state-of-the-art works in the field of PPML.
- Our construction obviates the need of a powerful machine on the client’s side as the classification takes place on the CSP. Hence, the role of client can be played by a constrained device whose sole responsibility will be to encrypt an outsource a file to the CSP.

B. Organization

The rest of the paper is organized as follows: In Section II, we present important published works in the area of PPML. In Section III, we show how to approximate the ReLU and Sigmoid activation functions using low degree Chebyshev polynomials. The methodology of our work is illustrated in Section IV, which is followed by extensive experimental results in Section V and finally, in Section VI we conclude the paper.

II. RELATED WORK

The first steps towards the development of a PPML model were implemented using Multiparty Computation (MPC). In an MPC scheme, the parties jointly compute a function while keeping the original inputs private. A plethora of different methods based on MPC has been proposed as a way to preserve the privacy of the underlying ML models [16]–[18].

Mohassel *et al.*, [19] designed SecureML – an efficient protocol for preserving the privacy of various ML models using MPC. This protocol is based on a two-server model, where the data owner first distributes the data among two *non-colluding* servers. These servers then train various models on the joint data using secure MPC. In addition, the protocol provides support to approximate the activation function during the training phase. Since SecureML requires changes in the training phase, the model does not apply to the problem of making the existing NN model oblivious. MiniONN [20] is another MPC-based approach used to convert any NN to an oblivious NN supporting privacy-preserving predictions with reasonable efficiency. Oblivious protocols are designed to

minimize the accuracy loss in linear transformation, pooling operations and activation functions. While MiniONN uses lightweight cryptographic primitives, such as garbled circuits and secret sharing, it still reveals information about the network (e.g. size of the filter) [21]. Wu *et al.* [22] proposed a privacy-preserving logistic regression model. As the logistic function is *not* linear, the authors use polynomial fitting to achieve a good approximation. However, by doing so the accuracy of the model is degraded. Graepel *et al.* [23] used Somewhat Homomorphic Encryption (SHE) [24] to train two simple classifiers: Linear Mean classifier [25] and Fisher’s Linear Discriminate classifiers [26] on encrypted data. To perform efficient computations on encrypted data, low degree polynomials were deployed. The proposed approach does not offer strong security guarantees as the model is leaked to the users.

Privacy-preserving classification for Naive Bayes, Hyperplane Decision, and Decision tree was provided by Bost *et al.* [27] using a combination of HE and garbled circuits. The basic idea was to provide users with the encrypted classification parameters and to compute the prediction values using a partially HE scheme [28]. However, this approach fails to preserve users’ privacy. This has been demonstrated in [29], where a *substitution-then-comparison* attack, is enough to reveal *all* information of one party to another.

A design based on Leveled HE (LHE) [30] is presented by Ehsan *et al.* [31]. Their goal was to utilize the properties of the LHE scheme to preserve the privacy of CNN while at the same time keep the accuracy as close as possible to the original model. In their work, the authors approximated the Sigmoid, ReLU and Tanh activation functions and achieved an accuracy of 99.52% on the MINST dataset [32]. This is a remarkable result, as the accuracy of the original model was measured at 99.56%. Unfortunately, their approach is computationally expensive, as both the training and testing phases are performed on encrypted data.

a) Most Relevant Related Work:: In [3], authors proposed CryptoNets – an NN model applied to encrypted data. While CryptoNets achieves remarkable accuracy, the construction is based on the use of square activation function. Hence, approximating a non-linear function causes instability during the training phase when the interval¹ is large. In our work, we overcome this problem by using Chebyshev approximation. More specifically, Chebyshev polynomials allows to approximate the activation function accurately even in larger intervals. Using Chebyshev polynomials, is *not* a novel idea. In [33], authors proposed a framework in which both the training and the classification phases are performed over encrypted data. In this work, authors used the Chebyshev polynomials to approximate the derivative of the activation functions. While their results are quite impressive, the classification process is computationally expensive. This is due to the fact that the ML model is encrypted and hence, classification is performed on encrypted data. To achieve better efficiency, we adapted a

¹By interval we mean the domain of definition of the activation function.

hybrid approach in which the client’s input is encrypted but the model is in plaintext.

III. PRELIMINARIES

A. Homomorphic Encryption

An HE scheme consists of the four following algorithms:

- KeyGen: A randomized algorithm that on input a security parameter λ , outputs a public/private key pair (pk, sk) .
- Enc: A randomized algorithm that on input a message m and a public key pk , outputs a ciphertext c .
- Dec: A deterministic algorithm that on input a ciphertext c and a secret key sk , outputs a plaintext m .
- Eval: A randomized algorithm that on input a public key pk , a function f , and a set of ciphertexts c_1, \dots, c_n , outputs $f(c_1, \dots, c_n)$.

B. Background on Polynomial Approximations

Approximating continuous functions is a problem that has drawn mathematicians’ attention for a very long time. While there are several ways to approximate a continuous function, in this work we are only interested in polynomial approximations. More specifically, we are using Chebyshev polynomials to approximate the Sigmoid and the ReLU functions. However, there are various works that use different approaches such as the x^2 function [3], the *Piecewise* approximation [14], lookup tables [12] etc. Unfortunately, all these methods face certain limitations. For example, the x^2 method can cause instability during the training phase and the creation of a piecewise linear approximation can sometimes be a complex optimization problem. With this in mind, we chose to work with Chebyshev polynomials.

$$T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x) \quad (1)$$

C. Chebyshev Polynomials

In this Section, we show how low degree Chebyshev polynomials can be utilized to approximate the activation functions. Chebyshev approximation is also known as the minimax approximation. The minimax polynomial approach is used for function approximation by improving the accuracy and lowering the overall computational complexity [34]. Instead of minimizing the error at the point of expansion like Taylor’s polynomial approximation, the minimax approach minimizes the error across a given input segment. The minimax approximation is used to find a mathematical function that minimizes the maximum error. As an example, for a function f defined over the interval $[a, b]$, the minimax approximation finds a polynomial $p(x)$ that minimizes $\max_{a \leq x \leq b} |f(x) - p(x)|$.

D. Chebyshev Approximation

To approximate a continuous function f , defined over $[a, b]$, we first need to express f as a series of Chebyshev polynomials at $[-1, 1]$. More precisely, f is expressed as: $f(x) = \sum_{k=0}^n c_k T_k(x)$, $x \in [-1, 1]$, where c_k is the Chebyshev coefficient and $T_k(x)$ can be calculated from equation 1. As a next step, we calculate the coefficients of the

polynomial and finally, express the polynomial in the original interval $[a, b]$.

Our results for approximating the Sigmoid and ReLU activation functions, using Chebyshev approximation, are illustrated in table I. The approximation error for both activation functions is calculated using equation $E(x) = f(x) - p(x)$

TABLE I
APPROXIMATING SIGMOID AND RELU

Approximation: Sigmoid				
x	Interval	Function(x)	Approximation	Difference
-4	[-5, 5]	0.017986	0.016360	-1.63e-03
-3	[-5, 5]	0.047426	0.049098	1.67e-03
-2	[-5, 5]	0.119203	0.118340	-8.63e-04
-1	[-5, 5]	0.268941	0.268522	-4.19e-04
1	[-5, 5]	0.731059	0.731478	4.19e-04
2	[-5, 5]	0.880797	0.881660	8.63e-04
3	[-5, 5]	0.952574	0.950902	-1.67e-03
4	[-5, 5]	0.982014	0.983640	1.63e-03
Approximation: ReLU				
-4	[-5, 5]	0.000000	-0.008871	-8.87e-03
-3	[-5, 5]	0.000000	0.014340	1.43e-02
-2	[-5, 5]	0.000000	-0.015085	-1.51e-02
-1	[-5, 5]	0.000000	-0.026883	-2.69e-02
1	[-5, 5]	1.000000	0.973117	-2.69e-02
2	[-5, 5]	2.000000	1.984915	-1.51e-02
3	[-5, 5]	3.000000	3.014340	1.43e-02
4	[-5, 5]	4.000000	3.991129	-8.87e-03

IV. METHODOLOGY

We start this section by describing our system model. The topology of our work is illustrated in figure 1.

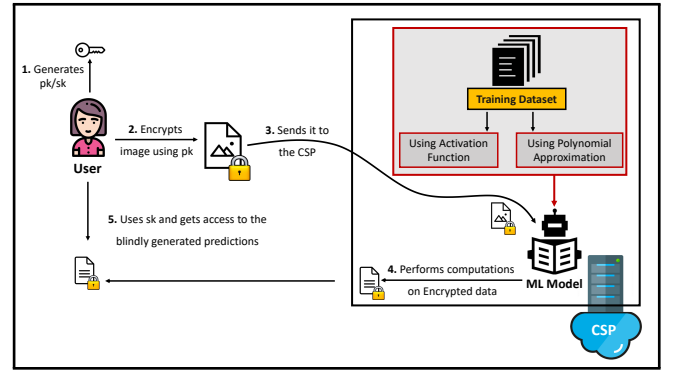


Fig. 1. Blind Faith High Level Overview

In our model, we consider a CNN capable of analyzing large volumes of data (images) in a variety of domains. The CNN is deployed in a privacy-preserving manner in the CSP. To preserve the users privacy, we use HE. Using an HE scheme allows us to perform computations on encrypted data. However, HE schemes face certain limitations as they only support addition and multiplication operations. Most of the operations in a CNN are simple additions and multiplications and can thus be evaluated using HE. However, activation functions are non linear and as a result, we cannot use HE to

- 1) Convolution: Input image 28×28 , Window size 5×5 , Stride(1,1), Number of input channels 1, number of output channels 5, Filters 5, Output $28 \times 28 \times 5$.
- 2) Activation Function: ReLU .
- 3) Pooling: Mean, Window size $2 \times 2 \times 1$, Stride(1,1), Output $14 \times 14 \times 5$.
- 4) Convolution: Window size 5×5 , Stride(1,1), Filters 10, Output $14 \times 14 \times 10$.
- 5) Activation Function: ReLU .
- 6) Pooling: Mean, Window size $2 \times 2 \times 1$, Stride(1,1), Output $7 \times 7 \times 10$.
- 7) Fully Connected: Fully connects the incoming $7 \times 7 \times 10$ nodes to the outgoing 128 nodes.
- 8) Activation: ReLU
- 9) Fully Connected: Fully connects the incoming 128 nodes to the outgoing 10 nodes.
- 10) Softmax: Generate a probability for 10 nodes

Fig. 2. Convolutional Neural Network for Training Phase

perform operations on them. To this end, we replace the activation functions with polynomial approximations as already discussed in section III-D. While higher degree polynomials would provide us with a better approximation, they would also introduce higher computation and communication costs and hence, render our construction inefficient.

Flow: The CNN model is deployed in a CSP and is trained using plaintext data. The weights and biases for this model are measured and made available to the CSP. For the training phase, we use the CNN given in figure 2. A user generates a public/private key pair for the HE scheme, encrypts an image and sends it to the CSP. Upon reception, the CSP runs the ML model and performs the classification in a privacy-preserving way.

A. Inference Phase

Although, the operations performed in the inference phase are nearly the same as in the training phase, there are few fundamental differences. For example, all operations in the inference phase are taking place on encrypted data while in contrast to the training phase where plaintext data is used. Similarly, the softmax which is part of the training phase is no longer available in the inference phase as shown in figure 3. This operation can be skipped because it is a monotonically increasing feature, and the cipher texts from the output of the last completely connected layer can be compared to give the prediction. Approximating or performing computation on encrypted data at this layer is out of the scope of the paper.

For the inference phase, we use the Fan-Vercauteren SHE scheme [24]. The reason for using this specific scheme is that it allows us to perform *both* addition and multiplication. Similarly, Single Instruction Multiple Data (SIMD) can be used to improve the performance of this scheme. Using SIMD technique, a batch of ciphertext is created in order to perform homomorphic evaluation in parallel. It is important to note

that this scheme has three important parameters that affect the security level, and its performance:

- Polynomial Modulus: This is an important parameter that affects the security level of the scheme. Polynomial modulus uses a power of two cyclotomic polynomial [35] and the recommended degrees for these polynomials are 1024, 2048, 4096, 8192 and beyond. On one side, a higher degree gives more security to the scheme while on the other side it degrades its performance.
- Coefficient Modulus: This parameter determines the Noise Budget (NB) in the encrypted ciphertext. The coefficient modulus is directly proportional to NB and inversely proportional to the security level of the scheme.
- Plaintext Modulus: The plaintext modulus affects NB in the freshly encrypted ciphertext. Additionally, it affects the NB consumption of homomorphic multiplications. For good performance, the recommendation is to keep the plaintext modulus as small as possible.

Each ciphertext in this encryption scheme has a specific quantity called NB – measured in bits. The NB is determined by the above parameters and consumed by the homomorphic operations. The consumption of the NB is based on the chosen encryption parameters. For addition operations, this budget consumption is almost negligible in comparison to multiplication operation. In sequential multiplication that occurs at the convolutional and fully connected layer, the consumption of NB is very high. Hence, it is important to reduce the multiplicative depth of the circuit by considering appropriate encryption parameters. Once the NB drops to zero, then the decryption of ciphertext is not possible. Therefore it is necessary to choose the parameters to be large enough to avoid this, but not so large that it becomes ineffective and non functional.

While the HE scheme is based on polynomials, user's input is provided as a real number. Therefore, there is a clear mismatch between the two. Hence, it is important to use an encoding scheme that maps one to the other. To this end, the user encodes the input using the plaintext modulus and then encrypts it using the public key. The user also generates the encryption parameter and shares it with the CSP. To perform computations on the encrypted data, the CSP must have access to these parameters.

Using the calculated weights and biases from the training phase and the encryption parameters, the CSP runs the inference phase on the encrypted image. The inference network is the same as the training network except that the activation functions are replaced by polynomial approximation and are built using an HE function.

Now, the activation functions are substituted by polynomials. Since these polynomials only have addition and multiplication operations that are supported by HE. Consequently, we can perform encrypted computations on these functions. Similarly, the pooling operation in the inference phase is straightforward – calculate the average value of four ciphertexts and multiply it with the appropriate values. However, the

- 1) *Convolution*: Input image 28×28 , Window size 5×5 , Stride(1,1), Number of input channels 1, number of output channels 5, Filters 5, Output $28 \times 28 \times 5$.
- 2) *Activation Function*: Approximated using polynomial approximation.
- 3) *Pooling*: Mean, Window size $2 \times 2 \times 1$, Stride(1,1), Output $14 \times 14 \times 5$.
- 4) *Convolution*: Window size 5×5 , Stride(1,1), Filters 10, Output $14 \times 14 \times 10$.
- 5) *Activation Function*: Approximated using polynomial approximation.
- 6) *Pooling*: Mean, Window size $2 \times 2 \times 1$, Stride(1,1), Output $7 \times 7 \times 10$.
- 7) *Fully Connected*: Fully connects the incoming $7 \times 7 \times 10$ nodes to the outgoing 128 nodes.
- 8) *Activation*: Approximated using polynomial approximation.
- 9) *Fully Connected*: Fully connects the incoming 128 nodes to the outgoing 10 nodes.

Fig. 3. Convolutional Neural Network for Inference Phase

convolutional layer is a bit more expensive in terms of NB as it is a sequence of multiplication operations.

Additionally, the softmax layer is not a part of the inference network. Using the inference network, the CSP performs computation on encrypted data and obtains an encrypted output. The CSP does not have access to the secret key and thus cannot access the result. Furthermore, as the softmax layer is removed from the inference network the CSP is not able to predict the final output of the layer.

At the end, the encrypted result of the output layer – an array of 10 values which are homomorphically encrypted – is sent back to the user. The user decrypts the results using the secret key and finds the output of the model which is the index corresponding to the highest among the 10 values.

At this point, it is important to highlight that the user utilized the ML model offered by the CSP and received the results without getting any valuable information about the underlying model. Similarly, the CSP ran the model on the encrypted image but at the same time was unable to extract any valuable information either for the content of the image or the actual prediction that sent back to the user. Hence, our model is considered as a privacy-preserving one.

V. PERFORMANCE EVALUATION

In this section, we present our experimental results.

a) *Experimental Setup*: All experiments were conducted in Python 3 using Ubuntu 18.04 LTS 64 bit (Intel Core i7, 2.80 GHz, 32GB). For the training phase, we used Tensor flow² to train our CNN model, while the actual experiments for that phase were conducted on Google Cola³ (with GPU enabled).

²<https://www.tensorflow.org/>

³<https://colab.research.google.com/>

Finally, for the inference phase we used Microsoft’s Simple Encrypted Arithmetic Library (SEAL) [36].

b) *Dataset*: To evaluate our model, similar to other works in the area, we used the MNIST dataset [32] which consists of 60,000 images of handwritten digits. To train our CNN model we used 50,000 images while the rest 10,000 were used for testing. Each image is 28×28 pixel array and is represented by its gray level in the range of 0-255.

A. Activation Function Approximation

As we mentioned in the previous sections, in our approach we use Chebyshev polynomials to approximate the ReLU and Sigmoid activation functions where inputs are images encrypted with an SHE scheme. The polynomial approximation of the ReLU activation function is shown in table II. Since the choices of the degree and the interval affect the performance of the model, it is necessary to choose suitable parameters. For this purpose, we conducted a series of experiments using different degrees and intervals⁴. As can be seen in table II, the activation functions are more accurately approximated when using polynomials of higher degree in small intervals. For example, the polynomial having degree 9 and interval $[-10, 10]$ more accurately approximate the ReLU function than the rest of the polynomials. The same applies to the Sigmoid activation function, where a high degree 9 and small interval $[-10, 10]$ give a better approximation as can be seen in III. However, the use of higher degree polynomials introduces a significant computation overhead, and small intervals limit the use of the approximation function.

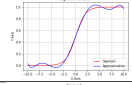
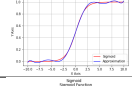
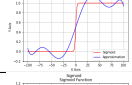
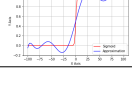
TABLE II
POLYNOMIAL APPROXIMATION OF THE RELU FUNCTION ON TWO INTERVALS $([-10,10], [-100,100])$ USING DIFFERENT DEGREES

Degree	Interval	Polynomial Approximation	ReLU Function
7	$[-10, 10]$	$(-8.88178419700125e-21) \times x^7 + (3.66197231323541e-6) \times x^6 + \dots$	
9	$[-10, 10]$	$(1.15960574476048e-21) \times x^9 - (7.031111115816643e-8) \times x^8 - \dots$	
7	$[-100, 100]$	$(-6.82121026329696e-27) \times x^7 + (3.6619723132354e-11) \times x^6 + \dots$	
9	$[-100, 100]$	$(1.12777343019843e-29) \times x^9 - (7.031111115816644e-15) \times x^8 - \dots$	

Furthermore, we performed a plethora of different experiments on the CNN model. We trained different networks by increasing the size of the convolutional layer and the size of the filters. We noticed that changing the number of convolutional layers and filters affects the overall accuracy of the network. As the size of the filter and layer increases, the accuracy of the network also increases. However, the efficiency of the

⁴Due to space constraints, the results for approximating the ReLU and Sigmoid using degrees 7, 9 are only included.

TABLE III
POLYNOMIAL APPROXIMATION OF THE SIGMOID FUNCTION ON TWO INTERVALS $[-10, 10]$, $[-100, 100]$ USING DIFFERENT DEGREES

Degree	Interval	Polynomial Approximation	Sigmoid Function
7	$[-10, 10]$	$(-4.34913635838155e - 7) \times x^7 - (5.82079696725621e - 16) \times x^6 + \dots$	
9	$[-10, 10]$	$(9.32721914680041e - 9) \times x^9 + (1.39698499452418e - 17) \times x^8 - \dots$	
7	$[-100, 100]$	$(-8.15672916212668e - 14) \times x^7 - (5.82076636528339e - 22) \times x^6 + \dots$	
9	$[-100, 100]$	$(2.59190909648308e - 17) \times x^9 + (1.3969822658824e - 25) \times x^8 - \dots$	

network drops significantly. Hence, for the training phase, we considered the network given in figure 2. First, we trained the CNN model using the ReLU activation function. The measured accuracy for that part was 99.2%. Then the same network was trained using the polynomial approximation function where we got an accuracy of 98.5% – a result that is very close to the original activation function.

For comparison, we used the model proposed in CryptoNets [3] which is similar to the one proposed in our paper – a convolutional layer, fully connected layers and an average pooling layer as shown in table IV. Training the model with the ReLU activation function, the accuracy of our model was 99.20% whereas CryptoNets achieved a 99%. Similarly, for the approximated function we obtained an accuracy of 98.5% while CryptoNets achieved 98.95%. For the same network, the accuracy of the model proposed in [33] was 99.02% using the ReLU activation function and 99% using the approximated function.

B. Performance Analysis of Performing Computation on Encrypted Data

Now, we proceed by discussing how the use of homomorphic encryption can affect the performance of the NN model. In our work, we trained the CNN model on plaintext data while the classification was performed on the ciphertexts. As a result, we had to perform computations on two types of data – plaintext and ciphertext. For this purpose, we used the SEAL library⁵ that allowed us to perform computations on ciphertext. Although the use of SEAL is straightforward, we still had to define certain parameters (see Section IV-A).

We performed a series of experiments using different encryption parameters. First, we looked at the polynomial modulus – the encryption parameter used in SEAL. During the experiments we observed that a smaller value of polynomial modulus leads to a more efficient result but at the same time the accuracy is decreased. In contrast, a higher value of the polynomial modulus gives more accurate results, however,

⁵<https://github.com/Lab41/PySEAL>

degrades the performance. The second encryption parameter is the coefficient modulus that decides the NB in the freshly encrypted ciphertext. This parameter is automatically set by SEAL based on the value of security level and polynomial modulus. Finally, increasing the value of the plaintext modulus, decreases the consumption of the NB.

C. Comparison with the Existence Model

Finally, we compared our results with state-of-the-art privacy-preserving NNs that utilize HE. The work proposed in CryptoNets [3] is similar to ours. In CryptoNets, the model is trained on plaintext data and then the trained model is used for the classification of encrypted instances. In order to have a fair comparison, it is important to incorporate the same network used in both works. To this end, we used the CryptoNets model.

Instead of using the overall performance of the model we decided to equate each layer. As can be seen in table IV, our model outperforms CryptoNets at both the encryption and decryption times as well as in the activation layer.

TABLE IV
COMPARISON WITH THE PREVIOUS MODELS

Layer	Description	Time	
		CryptoNets	Blind Faith
Encryption	Encoding+Encryption	44.5	8.5242
1st convolution layer	Same except the value of stride	30	60.36
1st activation function	Activation function	81	6.62
1st pooling layer	Mean pooling operation	127	0.188
2nd convolutional layer	-	-	64.822
2nd activation Function	Activation function	10	0.199
2nd pooling layer	-	-	0.092
1st fully connected layer	Generates 10 output	1.6	12.1839
2nd fully connected layer	-	-	0.326
Decryption	Image decryption	3	0.0021

a) **Future Directions:** As a continuation of this work, we plan to explore the possibility of combining Functional Encryption (FE) [37], with HE in an attempt to drastically improve the efficiency without sacrificing the accuracy of our construction. More precisely, we believe that it is possible to encrypt a user’s data under an HE scheme, and then use a symmetric FE [38], [39] scheme in the inference phase of the CNN. Such an approach could significantly reduce the computational costs and the total running time of the algorithms.

b) **Open Science & Reproducible Research:** To support open science and reproducible research as well as to give the opportunity to other researchers to use, test and hopefully extend/enhance our model we plan to make all of our source code available both on github and Zenodo⁶.

⁶<https://zenodo.org/>

VI. POSSIBLE SOCIETAL IMPACT

Today's code created by machine learning is making surprisingly insightful moves, spotting previously undiscovered features in medical images, investing in shrewd trades on the stock market, composing music, writing articles, creating art, evaluating the effectiveness of teachers, deciding if an applicant is eligible for getting a loan and even choosing which students will be admitted into a college. Undoubtedly, ML models and their underlying applications are driving the big-data economy.

However, in practice, the systems using these models are incorporating proxies. It has been observed that in many cases these proxies are bound to be inexact and often unfair. Considering that big data is not going away and predictive models are, and increasingly will be, the tools we will rely on to run our institutions, deploy our resources and manage our lives, it is of paramount importance to bring transparency into the game. A possible solution would be to protect users' personal, in many cases unrelated to the application, data from a potentially unfair and biased ML model. Consequently, creating models capable of performing high accuracy predictions whilst being unable to learn anything about processed data or generated output could bring improved fairness to the age of data. With this work we hope to pave the way towards building digitally blind prediction/evaluation systems that will, by design, eliminate several proxies such as geography, gender, race, etc. and eventually have a tangible impact in building a fairer, democratic and unbiased societies.

REFERENCES

- [1] N. Islam, W. Puech, K. Hayat, and R. Brouzet, "Application of homomorphism to secure image sharing," *Optics Communications*, vol. 284, no. 19, pp. 4412–4429, 2011.
- [2] M. Kim and K. Lauter, "Private genome analysis through homomorphic encryption," in *BMC medical informatics and decision making*, vol. 15, p. S3, Springer, 2015.
- [3] R. Gilad-Bachrach, N. Dowlin, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, "Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy," in *International Conference on Machine Learning*, pp. 201–210, 2016.
- [4] N. Dowlin, R. Gilad-Bachrach, K. Laine, K. Lauter, M. Naehrig, and J. Wernsing, "Manual for using homomorphic encryption for bioinformatics," *Proceedings of the IEEE*, vol. 105, no. 3, 2017.
- [5] T. Shortell and A. Shokoufandeh, "Secure signal processing using fully homomorphic encryption," *IET Information Security*, 2019.
- [6] G. cloud, "Google prediction api." <https://cloud.google.com/prediction/>.
- [7] M. Azure, "Azure machine learning." <https://azure.microsoft.com/en-us/services/machine-learning/#features>.
- [8] E. Labs, "Ersatz labs." <http://www.ersatzlabs.com/>.
- [9] A. Michalas, N. Paladi, and C. Gehrman, "Security aspects of e-health systems migration to the cloud," in *e-Health Networking, Applications and Services (Healthcom), 2014 IEEE 16th International Conference on*, pp. 212–218, IEEE, 2014.
- [10] P. Thaine, "Perfectly privacy-preserving ai."
- [11] S. Garge and D. Galindo, "Neural networks for encrypted data using homomorphic encryption," 2018.
- [12] J. L. Crawford, C. Gentry, S. Halevi, D. Platt, and V. Shoup, "Doing real work with fhe: the case of logistic regression," in *Proceedings of the 6th Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, pp. 1–12, 2018.
- [13] E. Hesamifard, H. Takabi, and M. Ghasemi, "Cryptodl: towards deep learning over encrypted data," in *Annual Computer Security Applications Conference, Los Angeles, California, USA*, 2016.
- [14] H. Chabanne, A. de Wargny, J. Milgram, C. Morel, and E. Prouff, "Privacy-preserving classification on deep neural network," *IACR Cryptol. ePrint Arch.*, vol. 2017, p. 35, 2017.
- [15] E. Chou, J. Beal, D. Levy, S. Yeung, A. Haque, and L. Fei-Fei, "Faster cryptonets: Leveraging sparsity for real-world encrypted inference," *arXiv preprint arXiv:1811.09953*, 2018.
- [16] P. Bunn and R. Ostrovsky, "Secure two-party k-means clustering," in *Proceedings of the 14th ACM conference on Computer and communications security*, pp. 486–497, 2007.
- [17] Y. Lindell and B. Pinkas, "Privacy preserving data mining," in *Annual International Cryptology Conference*, Springer, 2000.
- [18] J. Vaidya, H. Yu, and X. Jiang, "Privacy-preserving svm classification," *Knowledge and Information Systems*, 2008.
- [19] P. Mohassel and Y. Zhang, "Secureml: A system for scalable privacy-preserving machine learning," in *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 19–38, IEEE, 2017.
- [20] J. Liu, M. Juuti, Y. Lu, and N. Asokan, "Oblivious neural network predictions via minionn transformations," in *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 619–631, 2017.
- [21] C. Juvekar, V. Vaikuntanathan, and A. Chandrakasan, "{GAZELLE}: A low latency framework for secure neural network inference," in *27th {USENIX} Security Symposium ({USENIX} Security 18)*, pp. 1651–1669, 2018.
- [22] S. Wu, T. Teruya, J. Kawamoto, J. Sakuma, and H. Kikuchi, "Privacy-preservation for stochastic gradient descent application to secure logistic regression," in *The 27th Annual Conference of the Japanese Society for Artificial Intelligence*, 2013.
- [23] T. Graepel, K. Lauter, and M. Naehrig, "MI confidential: Machine learning on encrypted data," in *International Conference on Information Security and Cryptology*, pp. 1–21, Springer, 2012.
- [24] J. Fan and F. Vercauteren, "Somewhat practical fully homomorphic encryption," *IACR Cryptol. ePrint Arch.*, 2012.
- [25] R. O. Duda, P. E. Hart, et al., *Pattern classification*. John Wiley & Sons, 2006.
- [26] R. A. Fisher, "The use of multiple measurements in taxonomic problems," *Annals of eugenics*, vol. 7, no. 2, pp. 179–188, 1936.
- [27] R. Bost, R. A. Popa, S. Tu, and S. Goldwasser, "Machine learning classification over encrypted data," in *NDSS*, 2015.
- [28] P. Paillier, "Public-key cryptosystems based on composite degree residuosity classes," in *International conference on the theory and applications of cryptographic techniques*, pp. 223–238, Springer, 1999.
- [29] C.-z. Gao, Q. Cheng, P. He, W. Susilo, and J. Li, "Privacy-preserving naive bayes classifiers secure against the substitution-then-comparison attack," *Information Sciences*, vol. 444, 2018.
- [30] Z. Brakerski, C. Gentry, and V. Vaikuntanathan, "(leveled) fully homomorphic encryption without bootstrapping," *ACM Transactions on Computation Theory (TOCT)*, vol. 6, no. 3, pp. 1–36, 2014.
- [31] E. Hesamifard, H. Takabi, and M. Ghasemi, "Cryptodl: Deep neural networks over encrypted data," *arXiv preprint arXiv:1711.05189*, 2017.
- [32] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 2010.
- [33] E. Hesamifard, H. Takabi, M. Ghasemi, and R. N. Wright, "Privacy-preserving machine learning as a service," *Proceedings on Privacy Enhancing Technologies*, vol. 2018, no. 3, pp. 123–142, 2018.
- [34] J. Schlessman, "Approximation of the sigmoid function and its derivative using a minimax approach," 2002.
- [35] R. Thangadurai, "On the coefficients of cyclotomic polynomials," *Cyclotomic fields and related topics (Pune, 1999)*, 2000.
- [36] "Microsoft SEAL (release 3.2)." <https://github.com/Microsoft/SEAL>, Feb. 2019. Microsoft Research, Redmond, WA.
- [37] D. Boneh, A. Sahai, and B. Waters, "Functional encryption: Definitions and challenges," in *Theory of Cryptography Conference*, pp. 253–273, Springer, 2011.
- [38] A. Bakas, A. Michalas, and A. Ullah, "(f) unctional sifting: A privacy-preserving reputation system through multi-input functional encryption," in *Secure IT Systems: 25th Nordic Conference, NordSec 2020, Virtual Event, November 23–24, 2020, Proceedings 25*, pp. 111–126, Springer, 2021.
- [39] A. Bakas and A. Michalas, "Multi-input functional encryption: efficient applications from symmetric primitives," in *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pp. 1105–1112, IEEE, 2020.