

Partially Oblivious Neural Network Inference

Panagiotis Rizomiliotis, Christos Diou, Aikaterini Triakosia, Ilias Kyrannas and Konstantinos Tserpes

Dep. of Informatics and Telematics, Harokopio University, Omirou 9, Athens, Greece

Keywords: Homomorphic Encryption, Neural Network.

Abstract: Oblivious inference is the task of outsourcing a ML model, like neural-networks, without disclosing critical and sensitive information, like the model's parameters. One of the most prominent solutions for secure oblivious inference is based on a powerful cryptographic tools, like Homomorphic Encryption (HE) and/or multi-party computation (MPC). Even though the implementation of oblivious inference systems schemes has impressively improved the last decade, there are still significant limitations on the ML models that they can practically implement. Especially when both the ML model and the input data's confidentiality must be protected. In this paper, we introduce the notion of partially oblivious inference. We empirically show that for neural network models, like CNNs, some information leakage can be acceptable. We therefore propose a novel trade-off between security and efficiency. In our research, we investigate the impact on security and inference runtime performance from the CNN model's weights partial leakage. We experimentally demonstrate that in a CIFAR-10 network we can leak up to 80% of the model's weights with practically no security impact, while the necessary HE-multiplications are performed four times faster.

1 INTRODUCTION

Artificial intelligence (AI), and in particular, machine learning (ML) technology is transforming almost every business in the world. ML provides the ability to obtain deep insights from data sets, and to create models that outperform any human or expert system in critical tasks, like face recognition, medical diagnosis and financial predictions. Many companies offer such ML-based operations as a service (Machine learning as a service, MLaaS). MLaaS facilitates clients to benefit from ML models without the cost of establishing and maintaining an inhouse ML system. There are three parties involved in the transaction; the data owner, the model owner and the infrastructure provider.

However, the use of ML models raises crucial security and privacy concerns. The data set used for the ML model training and/or the MLaaS client's input in the inference phase can leak sensitive personal or business information. To complete the scenery of security threats, in several applications, like medical or financial, the ML models are considered as the MLaaS provider's intellectual property, and they must be protected.

Oblivious inference is the task of running a ML model without disclosing the client's input, the

model's prediction and/or by protecting the ownership of the trained model. This field of research is also referred to as Privacy-preserved machine learning (PPML).

Several solutions for oblivious inference have been proposed that utilize powerful cryptographic tools, like Multi-party Computation (MPC) primitives and the Homomorphic Encryption (HE) schemes. MPC based protocols facilitate the computation of an arbitrary function on private input from multiple parties. These protocols have significant communication overhead as they are very interactive. On the other hand, using HE cryptography we are able to perform computations on encrypted data, but with significant computation and storage overhead.

Several PPML schemes have been proposed that are either based solely on one of these technologies or that they leverage a combination of them (hybrid schemes). So far, literature has focused on two attack models. It is assumed that either the model owner is also the infrastructure provider or that the ML model that it is used, it is publicly known. This is a reasonable choice, as in both cases, the ML model's weights can be used in plaintext form. That is that, the schemes designers avoid expensive computations between ciphertexts and thus, they introduce inference systems that are practical.

In this paper, we consider the use cases in which the ML model’s confidentiality must be protected. The service provider wants to outsource the ML prediction computation (for instance to a cloud provider or to an edge device). However, the ML model constitutes intellectual property and its privacy must be preserved.

Protecting both the client’s input data and the model’s privacy can increase prohibitively the computational complexity as all the computations must be performed between encrypted data. Just as a rough estimation, the runtime of a single HE multiplication increases ten times when it is performed between two ciphertexts compared to HE multiplication between a plaintext and a ciphertext. At the same time, HE multiplications between encrypted data (ciphertexts) increase significantly the accumulative level of noise and they limit the applicability of the HE schemes. Thus, they must be avoided when possible.

Building on this observation, we introduce the notion of *partially oblivious (PO) inference*. In a PO inference system, the ML model owner decides to leak some of the model’s weights in order to improve the efficiency of the inference process. PO inference can be seen as a generalization of oblivious inference that offers a trade-off between security and efficiency. The PO inference systems lie between the two extreme use cases, the most secure but the less efficient in which all the ML model weights are encrypted and the less secure and the most efficient in which all the weights are revealed. The optimal point of equilibrium between efficiency and security depends on the use case.

Our work is summarized as follows:

1. We introduce the notion of Partially Oblivious inference for ML models.
2. We provide a security definition for the evaluation of the information leakage impact. In our analysis, the attacker is passive (“honest-but-curious”) and she aims to compute a model that simulates the protected one as accurately as possible. We use accuracy improvement as our security metric.
3. As a proof-of-concept use case, we apply the notion of PO inference to protect Convolutional Neural Networks (CNN) inference.
4. We experimentally measure the security and performance trade-off. We use two models trained with the MNIST (LeCun et al.,) and CIFAR-10 datasets (Krizhevsky, 2009), respectively. For the PO inference implementation, Gazelle-like (Juvekar et al., 2018) approach is used. Impressively, it is shown that in some scenarios, leakage of more than 80% of the model weight’s can be acceptable.

The paper is organized as follows. In Section 2,

the necessary background is provided. In Section 3, we analyze our motivation, we introduce the security attack model and the security definition for PO inference and we demonstrate the application of the PO inference to CNN models. Finally, in Section 4, we implement and evaluate the two CNN models and in Section 5, we conclude the paper.

1.1 Related Work

There are several lines of work for PPML systems that leverage advanced cryptographic tools, like MPC and HE. The most promising solutions are hybrid, and they are using HE to protect the linear and MPC to protect the non-linear layers.

CryptoNet ((Gilad-Bachrach et al., 2016)) is the first scheme that deploys the HE primitive for PPML on the MNIST benchmark. In the same research line, CHET (Dathathri et al., 2019b), SEALion (van Elsloo et al., 2019) and Faster Cryptonets (Chou et al., 2018) use HE and retrained networks. There are HE based schemes that use pre-trained networks, like Chimera (Boura et al., 2020) and Pegasus (Lu et al., 2020). In the pre-trained PPML category, we can find several proposals that use only MPC schemes, like ABY3 (Mohassel and Rindal, 2018), and XONN (Rhazi et al., 2019).

The most promising type of PPML systems are hybrid, i.e. the proposals that use both MPC and the HE schemes. Hybrid HE-MPC schemes provide an elegant solution for pre-trained networks. The MPC is responsible for the non-linear part (activation function) and HE for the linear transformations (FC and convolutional layers). Gazelle (Juvekar et al., 2018) is a state-of-the-art hybrid scheme for CNN prediction and several works have followed, like Delphi (Mishra et al., 2020), nGraph-HE (Boemer et al., 2019b), nGraph-HE2 (Boemer et al., 2019a), and PlainML-HE (Chen et al., 2019). All these schemes assume that either the model owner runs the models locally or that the ML model is publicly known.

There are several open source HE libraries that implement the operations of a HE schemes and offer higher-level API (Viand et al., 2021) and there is an ongoing effort to standardize APIs for HE schemes (sta, 2018). However, dealing directly with HE-libraries and operations is still a very challenging task for the developers. In order to facilitate developers work, HE compilers have been proposed to offer a high-level abstraction. There is a nice overview of existing HE-compilers in (Viand et al., 2021).

2 BACKGROUND

2.1 Homomorphic Encryption

In the last decade, the performance of HE schemes has impressively improved up to several orders of magnitude thanks to advances in the theory and to more efficient implementations. However, it is still significantly slower than plaintext computations, while realizing HE-based computations is complex for the non-expert.

Modern HE schemes belong into one of two main categories. The schemes that compute logical gates and thus, they are most efficient for generic applications, and the schemes that operate nicely on arithmetic or p -ary circuits and thus, they are used for the evaluation of polynomial functions. The CKKS (Cheon et al., 2017) scheme belongs in the second category. As it operates to arithmetic circuits on complex and real approximate numbers, CKKS is suitable for machine learning applications. We are going to use it in our experiments.

Following the last version of the HE Standard (sta, 2018), all the schemes must support the following types of operations: key and parameters management, encryption and decryption operation, HE evaluation of additions and multiplications, and noise management.

2.2 HE Evaluation Operations Cost

Practically, all the modern HE schemes are based on the hardness of the *Learning With Errors (LWE)* problem (Regev, 2005) and its polynomial ring variant. Depending on the scheme the plaintext, the keys, and the ciphertexts are elements of \mathbf{Z}_q^n or $\mathbf{Z}_q[X]/(X^n + 1)$, i.e. they are either vectors of integers or polynomials with integer coefficients.

In order to protect a message m a randomly selected vector (or polynomial) e is selected from a distribution and it is added to produce a noisy version of m . The level B of this added noise must always be between two bounds B_{min} and B_{max} . When $B < B_{min}$, the ciphertext cannot protect the message, while when $B > B_{max}$, the noise cannot be removed and the correct message cannot be retrieved anymore.

Thus, it is crucial to manage the level of noise induced by the HE operations. It has been demonstrated that the best noise management approach is to treat the ciphertext's noise level B as an invariant. That is that, after each HE operation, the level of noise must remain close to B .

In the CKKS (Cheon et al., 2017) scheme, the ciphertext is a pair of polynomials $c = (c_0, c_1)$ over

a ring of polynomials $\mathbf{Z}_q[X]/(X^N + 1)$, for appropriately selected integers q and N . The four main evaluation operations of CKKS scheme are summarized as:

1. Plaintext-Ciphertext Addition.

Let m and m' be two plaintexts and $c' = (c'_0, c'_1)$ be the encrypted value of m' . The output of the addition is $c_{output} = (m + c'_0, c'_1)$ and decrypts to $m + m'$ and the noise level is B .

2. Ciphertext Addition.

Let $c = (c_0, c_1)$ and $c' = (c'_0, c'_1)$ be the encrypted values of plaintexts m and m' . The output of the addition is $c_{output} = (c_0 + c'_0, c_1 + c'_1)$ and it is the ciphertext of $m + m'$ (approximately with good accuracy). The level of noise is upper bounded by $2B$.

3. Plaintext-Ciphertext Multiplication.

Let m and m' be two plaintexts and $c' = (c'_0, c'_1)$ be the encrypted value of m' . The output $c_{output} = (m \cdot c'_0, m \cdot c'_1)$ decrypts to $m \cdot m'$ and the level of noise is mB .

4. Ciphertext Multiplication.

Let $c = (c_0, c_1)$ and $c' = (c'_0, c'_1)$ be the encrypted values of plaintexts m and m' . The output of the multiplication is three polynomials, $c_{output} = (c_0 \cdot c'_0, c_0 \cdot c'_1 + c'_0 \cdot c_1, c_1 \cdot c'_1)$ and the noise level is B^2 .

It is clear that the ciphertext multiplication is the problematic one. The number of ciphertext polynomials increases linearly (one more polynomial after each multiplication) and the noise level increases exponentially (it becomes B^{2^L} , after L consecutive multiplications). To manage this size and noise increase, two refresh type operations are applied. In order to bring the dimension of the output ciphertext back to two, the relinearization algorithm is used. The resulting ciphertext c''_{output} is an encryption of the $m \cdot m'$ and the level of noise is B^2 . For the noise management, an algorithm called rescale (or modulo switching in other HE schemes) is used. However, it can be applied only a limited and predetermined number of times, usually equal to the multiplicative depth L of the arithmetic circuit.

Both algorithms, rescaling and relinearization, are costly in terms of computational complexity and both of them are applied after each multiplication between two ciphertexts. Relinearization has approximately the same computational cost with ciphertext multiplication and an evaluation key is required. The evaluations keys are created by the encryptor and passed to the evaluator.

To summarize, the HE multiplication between ciphertexts is a very costly operation in terms of com-

putational overhead and noise management. Compared to ciphertext multiplication, the other three HE evaluation operations are practically for free.

2.3 Plaintext Packing

One of the main features of some HE schemes that extremely improve performance is plaintext packing (also referred to as batching). It allows several scalar values to be encoded in the same plaintext. Thus, for schemes with cyclotomic polynomial of degree N , it is possible to store up to $N/2$ values in the same plaintext (we refer to them as slots). Thus, homomorphic operations can be performed component-wise in Single Instruction Multiple Data (SIMD) manner. This encoding has several limitations, since there is not random access operation and only cyclic rotations of the slots is allowed.

There are various choices for plaintext packing in ML, i.e. how the input data and the model weights are organized in plaintexts (or ciphertexts). Depending on the workload two are the main packing approaches, *batch-axis-packing* and *inter-axis packing*.

The *batch-axis-packing* is used by CryptoNets, nGraph-HE and nGraph-HE2. It is used to a 4D tensor of shape (B, C, H, W) , where B is the batch size, C is number of channels and H, W the height and width of input, along the batch axis. That is that, each plaintext (or ciphertext) packing has B slots and $C \cdot H \cdot W$ are needed. This approach assumes that B inputs are available for each inference operation.

On the other hand, *inter-axis packing* is used when each input is processed separately, i.e. it is not necessary to collect B inputs before performing a prediction (this is common in medical diagnosis). There are several packing choices, all of them encode scalars from the same input. This approach is used by Gazelle in which different packing is used for each type of linear transformation. We will use *inter-axis packing* in our analysis. In Section 3.4, we provide more details on the different inter-axis packing choices.

2.4 CNN Models

The neural-network inference has been identified as the main application area for privacy preserving technology, and especially for HE and MPC schemes, as we have seen in Section 1.1. However, there are practical limits to the complexity of the use cases that can be implemented (the unprotected computation must be at most a few hundreds of milliseconds).

A CNN model consists of linear layers (like convolutional layer and fully connected layer) and nonlinear layers, like an activation function, usually a

ReLU functions or a pooling function, like max-pooling. A very simple CNN appears in Fig. 1.

A fully connected (FC) layer with M inputs and N Outputs is specified by a tuple (\mathbf{W}, \mathbf{b}) , where \mathbf{W} is an $M \times N$ weight matrix and \mathbf{b} is a vector of length N , called bias. This layer receives as input vector \mathbf{v}_{in} of length M and computes the output as the linear transformation of the input: $\mathbf{v}_{out} = \mathbf{W}^T \mathbf{v}_{in} + \mathbf{b}$.

The convolutional layer has c_i number of input channels with image dimension $w_i \times h_i$ each and produces c_o output images with dimension $w_o \times h_o$. The Conv layer is parameterized by $c_i \times c_o$ many $f_w \times f_h$ filters.

3 PARTIALLY OBLIVIOUS INFERENCE

3.1 Attack Model

The architecture of the neural network (number of layers, type of neural network) is publicly known. On the other hand, the network's weights constitute intellectual property of MLaaS provider and their confidentiality must be protected. We make no assumptions regarding the training data or the training process. The training can be based on a private dataset as well as on a partially public one. Also, the ML model may have been trained from scratch or it may have been based on publicly known pre-trained model.

All the model inference computations are outsourced to a cloud provider or an edge device (we refer to both as the Server). We assume that the Server is *honest-but-curious*, i.e. it executes the operations correctly, but it wants to reveal any information that it can. The goal of the attacker is to compute a ML model that can simulate the original one as accurately as possible.

Our scenario appears in Fig. 2. We make the assumption that HE schemes are used at least for the linear layers of the ML model. In our experiments, we assume that the nonlinear layers are implemented using an MPC protocol (in our case the Garbled Circuit from the Gazelle system). However, our attack model is general and it can easily be adapted to other privacy preserving technologies as well.

The model's owner together with the data owner produce the necessary HE keys, namely, the secret key for the decryption, the public key for the encryption and the corresponding evaluation keys that are sent to the server. In a simplified version of this scenario, the model's and data owner is the same entity. The model's weights and the input data are encrypted with the public key and they are sent to the server.



Figure 1: The example of a simple CNN. The green blocks are the linear layers (2 convolutional and one fully connected) and blue blocks are the non-linear layers.

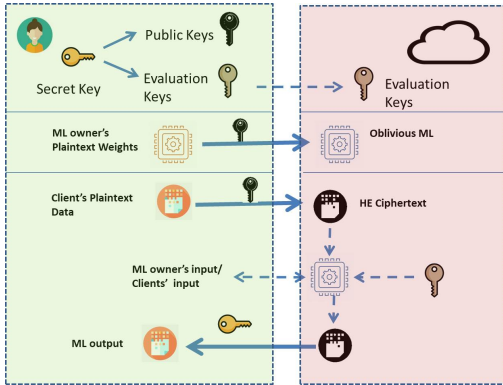


Figure 2: Oblivious Inference attack model.

The server runs the encrypted model on the encrypted input data, using the evaluation keys for the HE protected linear layers and any other technology for the nonlinear ones. The produced output is computed encrypted and it is sent to the legitimate user (it depends on the use case) who can decrypt it using the secret key directly or an MPC protocol on secret key's shares.

Note that, in theory, we can even hide the architecture of the model, however this is prohibitively expensive and it is avoided in practice.

3.2 Motivation

The efficiency, and practicality, of a HE-based system depends strongly on the type of operations that are performed. Based on the analysis in Section 2.2, two parameters are crucial:

1. *Multiplicative Depth*: The number of consecutive ciphertext multiplications must remain small. As the depth L increases, the modulo coefficient q must also increase and the HE scheme becomes inefficient.
2. *Number of ciphertext multiplications*: Even when multiplications can be performed in parallel, the computational complexity is significant. The goal is to reduce the number of ciphertext multiplications in total.

In the use cases under consideration the number of ciphertext multiplications is very high since both the system's input values and the ML coefficients are HE-encrypted.

Table 1: CKKS operations for security level $\lambda = 128$ -bit. N is the degree of the polynomial ($N/2$ slots per plaintext/ciphertext) and L is the multiplicative depth. All the parameters are defined in SEAL from the HE standard (sta, 2018).

	$N = 2^{12}$	$N = 2^{13}$	$N = 2^{14}$
	$L = 2$	$L = 4$	$L = 8$
Plain Mult.	1	4	16
Ciph Mult.	2.7	12.3	60
Rescale	7.6	38.5	175.7
Relinearization	16.1	80	477

One of the main techniques used to reduce the number of the multiplications between ciphertexts is packing (see Section 2.3), i.e. to organize several data in the same ciphertext and to compute in parallel all the computations with a single HE ciphertext multiplication. Both the human expert and HE-compilers use a pre-processing phase aiming to optimize the use of packing. For instance CHET, introduced in (Dathathri et al., 2019a), is a compiler that leverages the huge batching capacity and the rotations of the CKKS scheme, to decrease the number of required ciphertext multiplications.

Our goal is to go beyond the capabilities of packing by building on the computational cost asymmetry between HE ciphertext and HE plaintext multiplications. In Table 1, we can see an estimation of the runtime cost for different HE-multiplication related operations, i.e. plaintext multiplication, ciphertext multiplication, rescale and relinearization. The CKKS-RNS is used implemented in the SEAL library.

Each entry of the Table 1 indicates how many times lower is each operation compared to the plaintext multiplication for $N = 2^{12}$ and multiplicative depth $L = 2$. For instance for the same parameters N and L the ciphertext multiplication is 2.7 times slower and the relinearization is 16 times slower.

While we evaluate each operation separately, in practice the ciphertext multiplication is always followed by the relinearization operation. That it that, it is almost 19 times slower in practice. Similarly, we can argue for the rescale operation.

From Table 1, it is clear that the multiplications between a plaintext and ciphertext are much more efficient. Based on this observation, we are motivated to investigate the possibility of leaking information that has limited impact on the security of the protected scheme. The performance benefits are pretty clear in

terms of runtime overhead as the revealed values can be used in plaintext form.

Note. We need to have in mind that plaintext multiplication form leads also to limited noise growth and as a result a larger variety of functions that can be HE implemented. This is illustrate in Example 1. In this paper, we don't investigate this positive side-effect.

Example 1. Let $f(x, y, z, w) = xy + zw$ be a bilinear function that must be computed homomorphically. The naive approach requires two ciphertext multiplications, and a single ciphertext addition.

Let's assume that any two of the four input values can be leaked, i.e. they can be used as plaintext. There are $\binom{4}{2}$ different combinations of two inputs, i.e. 6 in total. Due to the symmetry of the function they belong into two equivalent classes, either both inputs are the operands of the same multiplication (i.e. $\{x, y\}$ or $\{z, w\}$) or operands of different multiplications.

In the first case, it requires one ciphertext multiplication, and one plaintext-ciphertext addition (the multiplication between plaintexts is for free). In the second case, it requires two plaintext-ciphertext multiplications, and one ciphertext addition. This is much more efficient.

3.3 Security Definition

In this section, we introduce the notion of *Partially Oblivious (PO)* inference and we provide the corresponding security definition.

Let \mathcal{W} be the set of weights of a model \mathcal{M} and let \mathcal{L} be a subset of \mathcal{W} .

Definition 1. A model \mathcal{M} is \mathcal{L} -Partially Oblivious (\mathcal{L} -PO), when the model is FHE computed and only the weights $w \notin \mathcal{L}$ are HE-encrypted.

The definition implies that the model's inference is computed using HE while the weights that are in the subset \mathcal{L} are used unencrypted, i.e. these weights are leaked.

When $\mathcal{L} = \emptyset$, all the weights are encrypted and we have the standard definition of oblivious inference. The model is \emptyset -PO.

Next, we introduce a security definition to assess the impact of the proposed information leakage. In our attack scenario, we assume that the adversary's goal is to steal the model. That is that, the attacker wants to produce a model $\hat{\mathcal{M}}$ that is equivalent to \mathcal{M} .

Let \mathcal{M}_0 and $\hat{\mathcal{M}}_{\mathcal{L}}$ be the models that the attacker computes when all the weights are encrypted and when \mathcal{L} are leaked, respectively. Also, let ACC_0 and $ACC_{\mathcal{L}}$ be the accuracy of each model.

Definition 2. A \mathcal{L} -PO model \mathcal{M} is λ -secure, if the advantage of any polynomial-time adversary A ,

$$Adv_A(\mathcal{M}, \mathcal{L}) = ACC_{\mathcal{L}} - ACC_0 \quad (1)$$

is upper bounded by λ , i.e.

$$Adv_A(\mathcal{M}, \mathcal{L}) \leq \lambda.$$

Ideally, $\lambda = 0$. In that case, the leaked weights \mathcal{L} do not improve the attackers capability to steal the model.

3.4 CNN Mdel Partially Oblivious Inference

In this section, we investigate the application of our Partially Oblivious inference approach to the CNN use case. From our attack model, the topology of the network is publicly known, but the weights are confidential.

Our goal is to investigate trade-offs between security and performance. This is expressed as information leakage and more precisely as revealing model's weights. We will show that the model's owner can use some of the weights in plaintext in order to improve inference runtime performance, while at the same time this leakage gives a very limited advantage to the attacker. These weights are only used in the linear layers.

The linear layers (linear operations of fully connected and convolutional layers) are implemented using an HE scheme. To simplify our analysis, we assume that the non-linear layers (activation functions) refresh the ciphertext noise. This is very common in the hybrid schemes. For instance in Gazelle (Juvekar et al., 2018), the non-linear layers are implemented with MPC (using Garbled Circuits) and the output of the layer is a ciphertext of with fresh noise (the same applies with TEE based solutions).

The model owner can reveal a certain percentage of the model's weights. The higher the percentage the more efficient the inference process. At the same time the security level is decreasing. The selection of the weights is subject to the restrictions imposed by the packing policy. The weights that are encoded in the same polynomial must be treated as a group. That is that, either they will all be revealed and used in plaintext form or they must all be protected and used in ciphertext form.

In our investigation, we follow the inter-axis packing (see Section 2.3) which is more common in medical diagnosis use cases and we analyze different packing techniques for the convolutional and fully connected layers. For the fully connected layer, we analyze the three packing techniques from Gazelle (Juvekar et al., 2018), namely the naive approach, the

diagonal and the hybrid. On the other hand, the two packing techniques for the convolutional layer (called in the paper padded and packed SISO) treat each filter value independently. Thus, the model owner can decide on the confidentiality of the convolutional layer weights without restrictions from the packing techniques used.

The model owner selects a percentage of the model weights to reveal following one of the following strategies:

1. *Random selection*: The weight groups of the fully connected layer and the individual filter values of the convolutional layer are selected completely at random.
2. *Maximum weight*: The weight groups of the fully connected layer and the individual filter weights of the convolutional layer with the largest mean of absolute values are selected.

For example, consider a fully connected layer with M inputs and N outputs. This layer is represented by an $M \times N$ weight matrix \mathbf{W} . The linear output (logit) is $\mathbf{z} = \mathbf{W}^T \mathbf{h}_{in} + \mathbf{b}$, where \mathbf{h}_{in} is the input vector and \mathbf{b} the biases. The naive approach with rows as groups will select $\lfloor pM \rfloor$ rows of \mathbf{W} for encryption. The first policy will randomly select the rows, while the second will select rows i with the highest mean $\frac{1}{M \cdot N} \sum_j |w_{ij}|$.

In the case of a convolutional layer, the layer is represented by a $k \times k \times M \times N$ tensor \mathbf{W} , where k is the kernel size, M is the number of input channels and N is the number of output channels. In this case, a group can be a “filter” W_i with dimensions $k \times k \times M$, or each $k \times k$ kernel W_{ij} . For the maximum weight strategy, the filters with the largest mean of absolute weights are selected.

For both FC and convolutional layers, encrypting the biases \mathbf{b} is an additional option (the biases are also encrypted as a group).

Regarding the attacker, we assume that she is trying to produce a model from the leaked information. We assume that the attacker has a very small set of data just for the evaluation, but not sufficient data to train or fine-tune a model. The attacker follows one of the following policies for the prediction of the missing weights.

1. *Constant (0.0)*: All the weights are replaced by the constant value (zero, in the case of our experiments).
2. *Mean (μ)*: The mean value of the known weights of the same layer is used. If no weights of the current layer are known, then the constant policy is used for that layer.
3. *Normal ($\mathcal{N}(0,1)$)*: The values are sampled from a standard normal distribution.

Table 2: Neural network architectures used in the experiments with MNIST and CIFAR-10 datasets. “conv_3 × 3_X” denotes a convolutional layer with X filters and a 3 × 3 kernel. FC-Y is a fully connected layer with Y neurons, while “batchnorm” denotes batch normalization (Ioffe and Szegedy, 2015). The ReLU activation function is used in all layers except the output FC layer.

model for MNIST	model for CIFAR-10
Input: (28 × 28 × 1)	Input: (32 × 32 × 3)
conv_3 × 3_32	conv_3 × 3_32
maxpool_2 × 2	batchnorm
conv_3 × 3_64	conv_3 × 3_32
maxpool_2 × 2	batchnorm
FC-10	maxpool_2 × 2
softmax	conv_3 × 3_64
	batchnorm
	conv_3 × 3_64
	batchnorm
	maxpool_2 × 2
	conv_3 × 3_128
	batchnorm
	conv_3 × 3_128
	batchnorm
	maxpool_2 × 2
	FC-128
	Dropout ($p = 0.5$)
	FC-10
	softmax

4. *Fitted normal ($\mathcal{N}(\mu, \sigma)$)*: Same as the *Normal* policy, but the values are sampled from a normal distribution that is estimated from the known weights of the same layer (i.e., with mean equal to the known weight mean, and standard deviation equal to the unbiased estimate of the standard deviation of the known weights).

In the following section, we empirically evaluate the tradeoffs between computational efficiency and security of the proposed \mathcal{L} -PO inference.

4 EXPERIMENTS

4.1 Experimental Setup

For our experiments we have used the MNIST (LeCun et al.,) and CIFAR-10 (Krizhevsky, 2009) data sets, using the standard train/test splits. The architecture of the networks used in each dataset are outlined in Table 2.

These architectures are used for empirical evaluation of the security and computational efficiency of the different \mathcal{L} -PO strategies.

4.2 \mathcal{L} -PO Security

To evaluate the security of \mathcal{L} -PO for CNN networks we initially train both network architectures of Table 2 using the Adam optimizer (Kingma and Ba, 2014) with a learning rate of 0.001 and categorical cross-entropy loss. The MNIST network is trained for 10 epochs without the use of a validation set, while the CIFAR-10 network is trained for 20 epochs, with 20% validation set and early stopping if no reduction in loss is observed for more than 3 epochs.

Then, we apply an “encoding” step, where weights of the neural network are selected according to the strategies described in Section 3.4. In this step we simulate \mathcal{L} -PO by storing the indices of the weights that would be selected for encryption. Then, we evaluate the accuracy that an attacker would achieve in a “decoding” step, for different policies (also described in Section 3.4).

Both the original model \mathcal{M} , the model $\hat{\mathcal{M}}$ that is estimated by the attacker, as well as the fully encrypted model \mathcal{M}_0 are evaluated in the test set of each dataset and the resulting accuracies are used to assess the security of \mathcal{L} -PO in terms of $Adv_A(\mathcal{M}, \mathcal{L})$ defined in Section 3.3.

In each experiment we define the percentage p of the weights to be selected for encryption, the weight selection strategy, whether to select biases for encryption, as well as the attacker policy. Since some of the policies applied by the attacker are stochastic, we repeat each experiment 10 times and report the average and standard deviation of $Adv_A(\mathcal{M}, \mathcal{L})$. For each experiment (i.e., combination of these options), the original network is trained only once and is used in all 10 iterations. Training a network for each different experiment (instead of using the same network across all experiments of the same datasets) helps take into account randomness introduced by model training (e.g., due to weight initialization).

Results of the experiments are shown in Table 3, for both datasets and for some combinations of p , weight selection strategies and attacker prediction policies. For each experiment, the table provides the average and standard deviation of the attacker advantage observed in the 10 experiment runs. In all experiments, bias weights have been selected for encryption.

Note also that, after encrypting all weights ($p = 1.0$) the model accuracy is roughly equal to the class prior (i.e., the accuracy of the random classifier). This ascertains our knowledge of the model architecture by itself does not lead to any attacker advantage.

An important observation from these results is that for both datasets and models, the attacker advantage

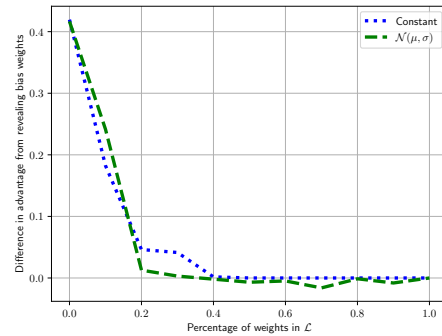


Figure 3: Difference in $Adv_A(\mathcal{M}, \mathcal{L})$ achieved by an attacker by when biases are not included in \mathcal{L} , for the max weight selection strategy. For very small p , the attacker can benefit from observing the biases, however as p grows this advantage quickly becomes insignificant.

quickly diminishes as more weights are encrypted. This is more pronounced in the more complex model of the CIFAR-10 dataset, where for $p > 0.2$, the attacker advantage is zero for both weight selection strategies. But even for a small number of hidden weights, e.g., $p = 0.1$, the maximum attacker advantage (obtained with the constant policy) is only 47.9, as opposed to 70.5 if she or he had access to the full model. In this case, the attacker would achieve approximately 0.58 accuracy, while with the full model she would achieve approximately 0.8. For the max weight selection strategy the advantage is even lower.

Another interesting observation is that the weight selection strategy plays an important role especially for the CIFAR-10 dataset. Max weight selection seems to be the most effective for both datasets. On the other hand, random filter and random weight selection conveys the minimum information about how weights were selected to an adversary. It seems that the optimal weight selection strategy depends on the model and more sophisticated methods could be explored. This is not further discussed in this paper and is left as future work. In the ideal case, one should evaluate different weight selection strategies and use the one that seems to provide the best results for each model.

Regarding the different attacker weight estimation policies, replacing all weights with zero leads to very good results for both datasets, while using weights from a fitted normal distribution seems to be a good policy as well. On the other hand, the mean and standard normal polices do not seem to be as effective.

In all experiments, the biases have been selected for encryption, since the addition is relatively cheap, computationally. Figure 3 illustrates the effect of including biases in \mathcal{L} for the CIFAR dataset for the random and the best two weight estimation policies for

Table 3: Results of experiments. Values are average and standard deviation across 10 runs, in $100Adv_A(\mathcal{M}, \mathcal{L})$ (a similar table with the corresponding accuracy values is provided in the appendix). All runs include hidden biases, while the asterisk ‘*’ indicates the model without any hidden weights. Columns $const = 0.0$, $\mathcal{N}(0, 1)$, $\mathcal{N}(\mu, \sigma)$ and μ correspond to the different weight estimation policies of the attacker.

Random weight selection								
p	MNIST				CIFAR10			
	$const = 0.0$	$\mathcal{N}(0, 1)$	$\mathcal{N}(\mu, \sigma)$	μ	$const = 0.0$	$\mathcal{N}(0, 1)$	$\mathcal{N}(\mu, \sigma)$	μ
*	89.2	88.7	87.8	89.2	70.5	70.9	70.5	70.5
0.0	89.1 (0.0)	27.4 (19.5)	86.6 (0.9)	89.1 (0.0)	47.9 (0.0)	3.1 (2.3)	39.9 (12.5)	40.6 (0.0)
0.1	88.3 (0.6)	0.7 (3.1)	82.2 (4.9)	88.4 (0.6)	33.3 (11.2)	0.0 (0.7)	11.6 (9.5)	18.4 (13.2)
0.2	85.6 (3.1)	0.6 (3.0)	62.2 (12.0)	83.8 (4.6)	18.2 (5.7)	0.8 (0.7)	5.2 (4.1)	9.2 (4.5)
0.3	83.8 (2.1)	0.7 (1.0)	47.0 (13.7)	83.1 (3.3)	0.4 (1.0)	0.7 (0.8)	1.2 (1.7)	0.5 (1.4)
0.4	78.3 (4.3)	0.1 (2.0)	29.8 (13.8)	73.4 (6.0)	0.6 (1.1)	0.9 (1.2)	-0.2 (1.4)	1.2 (1.5)
0.5	72.0 (9.0)	-0.8 (3.2)	19.1 (7.9)	67.9 (7.6)	-0.0 (0.5)	0.7 (1.5)	0.3 (1.9)	-0.0 (1.5)
0.6	47.6 (15.3)	0.3 (1.8)	8.8 (4.5)	40.5 (13.4)	-0.4 (0.6)	0.4 (0.3)	-0.3 (0.6)	0.0 (0.0)
0.7	28.3 (9.9)	0.6 (1.7)	-0.0 (4.3)	14.3 (3.8)	0.0 (0.0)	0.7 (1.1)	0.1 (1.1)	0.0 (0.0)
0.8	9.6 (7.2)	-0.8 (2.7)	2.0 (3.0)	9.6 (6.6)	-0.1 (0.4)	-0.1 (0.7)	-0.1 (0.7)	0.0 (0.0)
0.9	4.6 (5.9)	0.4 (3.3)	-1.2 (2.9)	-0.2 (1.5)	0.0 (0.0)	-0.2 (1.5)	0.6 (1.0)	0.0 (0.0)
1.0	0.0 (0.0)	0.0 (1.8)	0.0 (4.5)	0.0 (0.0)	0.0 (0.0)	0.0 (0.7)	0.0 (1.3)	0.0 (0.0)

Max weight selection								
p	MNIST				CIFAR10			
	$const = 0.0$	$\mathcal{N}(0, 1)$	$\mathcal{N}(\mu, \sigma)$	μ	$const = 0.0$	$\mathcal{N}(0, 1)$	$\mathcal{N}(\mu, \sigma)$	μ
*	89.0	87.8	89.3	89.0	69.7	69.7	69.9	69.7
0.0	88.8 (0.0)	16.5 (10.9)	86.7 (3.8)	88.7 (0.0)	28.0 (0.0)	2.5 (3.6)	27.7 (14.2)	17.9 (0.0)
0.1	75.5 (0.0)	0.2 (2.6)	71.0 (8.6)	79.9 (0.0)	22.2 (0.0)	0.1 (0.3)	13.4 (10.0)	19.8 (0.0)
0.2	77.9 (0.0)	-1.2 (2.6)	60.2 (10.0)	76.9 (0.0)	8.2 (0.0)	-0.2 (1.0)	3.5 (2.3)	8.9 (0.0)
0.3	67.4 (0.0)	-2.3 (2.6)	25.9 (14.7)	43.1 (0.0)	0.3 (0.0)	0.4 (0.4)	0.9 (1.8)	1.8 (0.0)
0.4	52.1 (0.0)	-0.5 (3.8)	14.3 (4.3)	24.1 (0.0)	0.0 (0.0)	0.0 (1.0)	0.2 (0.9)	0.6 (0.0)
0.5	33.4 (0.0)	-2.5 (2.3)	3.6 (3.3)	3.9 (0.0)	0.0 (0.0)	0.3 (0.8)	-0.2 (0.9)	0.0 (0.0)
0.6	39.4 (0.0)	-0.7 (2.9)	1.6 (2.7)	3.5 (0.0)	0.0 (0.0)	-0.2 (0.8)	0.4 (0.5)	1.4 (0.0)
0.7	15.9 (0.0)	-1.9 (1.2)	1.9 (3.1)	2.5 (0.0)	0.0 (0.0)	0.3 (0.8)	0.4 (1.1)	0.0 (0.0)
0.8	10.3 (0.0)	-0.6 (2.7)	1.1 (2.7)	8.1 (0.0)	0.0 (0.0)	-0.1 (0.7)	0.4 (0.9)	0.0 (0.0)
0.9	6.2 (0.0)	-0.4 (3.3)	1.0 (3.2)	6.9 (0.0)	0.0 (0.0)	-0.1 (0.3)	0.3 (1.0)	0.0 (0.0)
1.0	0.0 (0.0)	0.0 (2.3)	0.0 (2.4)	0.0 (0.0)	0.0 (0.0)	0.0 (0.6)	0.0 (1.3)	0.0 (0.0)

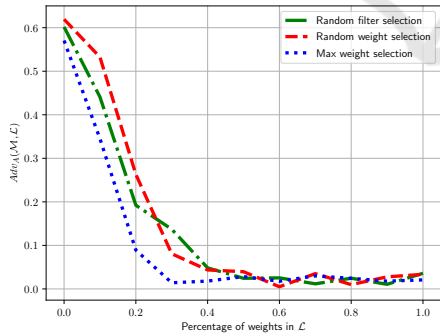


Figure 4: Worst-case across 10 runs, between random filter selection, random weight selection and max weight selection on the CIFAR-10 dataset, for varying percentages of the weights included in \mathcal{L} . In all cases, the reference accuracy is approximately $ACC_0 = 0.1$, (achieved when setting all weights and biases equal to zero). As p grows, the attacker advantage quickly diminishes.

that dataset.

Figure 4 illustrates the worst-case in the CIFAR-10 dataset. For each value of p , the plot shows the maximum advantage $Adv_A(\mathcal{M}, \mathcal{L})$ achieved by the at-

tacker across 10 runs and across all different weight estimation policies. Ever when encrypting a small percentage p of the weights, the model effectiveness drops significantly. For example, the maximum advantage for $p = 0.1$ and the max weight selection strategy is 0.34, leading to a model accuracy of approximately 0.44, which is significantly worse than the accuracy of the original model (close to 0.8). When encrypting more than 20% of the weights, the attacker advantage becomes insignificant, even in the worst case.

Overall, these results indicate that it is possible to only encrypt a portion of the weights of a neural network without enabling an attacker to infer the model weights. In addition, a significant drop in model effectiveness is observed even when hiding 10% of the weights. Finally, the security was higher for the more complex CIFAR-10 model, compared to the simpler MNIST model, especially for smaller percentages of hidden weights.

4.3 CNN Linear Layer Micro-benchmarks

In this section, we evaluate the impact that the selection of the model weights type (i.e. plaintext or ciphertext) has on the inference performance. We assume that a hybrid inference system is used and more precisely, a system similar to Gazelle.

As the model weights are only used in the linear layers, and in order to isolate the impact of the weights leakage to the system’s performance, we have implemented only the linear layers using the CKKS scheme from SEAL (SEAL, 2020). Regarding the nonlinear layers, we assume that the privacy preserving techniques, used to implement the nonlinear functions, refresh the noise induced by the HE operations of the linear layer. Such an implementation is used in Gazelle that leverages MPC protocols like Garbled Circuits. Thus, we can further isolate the impact of the leakage selection, as the multiplicative depth per linear layer is very small (usually $L = 1$ or $L = 2$) and we avoid expensive noise refreshing operations.

The linear layer weights (convolutional and fully connected) are homomorphically encrypted and different packing techniques can be used. For our analysis we use the packing techniques from Gazelle. For more details on these techniques please refer to the original paper (Juvekar et al., 2018).

All these techniques mainly perform ciphertext and/or plaintext HE multiplications between the layer input and the model weights. The resulting products are always ciphertexts. Depending on the packing technique, some additions and rotations on these ciphertexts are needed to produce the layer’s output. However, these last operations (additions and rotations) and their runtime overhead doesn’t depend on the initial type of the model weights, as they always operate on ciphertexts. That is that, only the initial HE multiplications performance reflects the weight’s leakage impact. Thus, we will only consider these multiplications in our experiments.

For all both the FC and the convolutional layers, the input of length is encrypted in a single ciphertext. In the FC layer with M inputs and N outputs, the matrix \mathbf{W} is partially encrypted up to percentage p . Depending on the packing used, the entries of \mathbf{W} are grouped together in ciphertexts or plaintexts. In the naive packing, each row constitutes a different group of M elements and there are N groups. Thus, we need $N \cdot (1 - p)$ plaintext multiplications and $N \cdot p$ ciphertext multiplications. In the diagonal packing, the number of multiplications depends on the inputs M . Similarly, we need $M \cdot (p - 1)$ plaintext and $M \cdot p$ ciphertext multiplications. Finally, the hybrid pack-

Table 4: Naive and diagonal packing for the FC layer. The multiplicative factor of runtime overhead compared to the all weights in plaintext computation.

p	(128, 10)	(2048, 128)	(128, 10)	(2048, 128)
0.0	1	1	1	1
0.1	2.9	4.4	2.8	4.7
0.2	4.8	7.9	4.7	8
0.3	6.7	11.6	6.6	11.8
0.4	8.7	15.1	8.8	15.3
0.5	10.4	18.3	10	18.5
0.6	12.5	22.1	12.3	22
0.7	14.3	25.3	14.2	25.5
0.8	15.9	29.2	16	30
0.9	18.2	32.5	18	32.8
1.0	19.8	35.7	20.1	36.2

ing needs the same number of multiplications as the naive techniques.

We evaluate the overhead of using a partially encrypted matrix \mathbf{W} in Table 4. The overhead is computed a multiplicative factor compared to the all weights in plaintext form case. Since after each linear layer the HE noise is refreshed, we avoid the rescale operation. Each ciphertext multiplication is followed by a relinearization operation.

In Table 5, we investigate the impact of the rescale operation. Since this operation is used after both the plaintext and the ciphertext operations, the relative performance gain from the weights leakage is reduced.

Finally, in the convolutional layer, each entry of the 3×3 filters are stored in a different ciphertext and almost all the slots are filled with the corresponding entry’s value. Thus, for each filter application we need 9 ciphertext or 9 plaintext multiplications. That is that, we assume that the whole filter is either encrypted or in plaintext form. Each convolutional layer input is and RGB image $32 \times 32 \times 3$. The result appears in Table 6.

In our experiments we use the convolutional and fully connected layers computed in Section 4.2 and we compute the performance of the linear layers for different values of the weights leakage percentage p .

5 CONCLUSIONS

This paper initiates a new line of research regarding the oblivious outsourcing of ML models computation. More specifically, we investigate the trade-offs between security and efficiency, when some information leakage is acceptable.

Our work serves mainly as a proof of concept using CNNs. We have shown that the model owner of a CIFAR-10 network can reveal 80% of selected model’s weights in order to reduce the linear lay-

Table 5: Naive packing for the FC layer using the rescale operation after the plaintext multiplication and using both the rescale and relinearization after the ciphertext multiplication. The multiplicative factor of runtime overhead compared to the all weights in plaintext computation.

p	(128, 10)	(2048, 128)
0.0	1	1
0.1	1.3	4.4
0.2	1.55	7.9
0.3	1.7	11.6
0.4	1.9	2.2
0.5	2.15	2.5
0.6	2.3	2.9
0.7	2.5	25.3
0.8	2.7	3.252
0.9	2.95	3.7
1.0	3.11	4.1

Table 6: The convolutional layer using the relinearization operation after the ciphertext multiplication. All the 9 multiplications for a filter application are either all ciphertext or all plaintext. The multiplicative factor of runtime overhead compared to the all weights in plaintext computation.

p	(3, 3)
0.0	1
0.1	4.5
0.2	8
0.3	11.7
0.4	15
0.5	18.15
0.6	22.3
0.7	25.5
0.8	29.7
0.9	32.55
1.0	36.11

ers cost of multiplications by a factor of 4. While similar security-performance trade-offs are very common in applied cryptography (in searchable symmetric schemes for instance), it is the first time that such approach is proposed in ML model inference.

Further research will follow. New attack models must be proposed and new more fine-grained security definitions must be introduced per use case. At the same time, the efficiency gain per use case must be evaluated both theoretically (complexity asymptotic) as well as experimentally. Our goal will be to leverage the results of this research and provide new design guidelines for efficient HE-compilers.

ACKNOWLEDGEMENTS

This work was supported by the project COLLABS, funded by the European Commission under Grant Agreements No. 871518. This publication reflects the views only of the authors, and the Commission cannot be held responsible for any use which may be made of the information contained therein.

REFERENCES

(2018). Homomorphic encryption standardization. <https://homomorphicencryption.org/standard/>.

Boemer, F., Costache, A., Cammarota, R., and Wierzynski, C. (2019a). ngraph-he2: A high-throughput framework for neural network inference on encrypted data. In Brenner, M., Lepoint, T., and Rohloff, K., editors, *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography, WAHC@CCS 2019, London, UK, November 11-15, 2019*, pages 45–56. ACM.

Boemer, F., Lao, Y., Cammarota, R., and Wierzynski, C. (2019b). ngraph-he: a graph compiler for deep learning on homomorphically encrypted data. In Palumbo, F., Becchi, M., Schulz, M., and Sato, K., editors, *Proceedings of the 16th ACM International Conference on Computing Frontiers, CF 2019, Alghero, Italy, April 30 - May 2, 2019*, pages 3–13. ACM.

Boura, C., Gama, N., Georgieva, M., and Jetchev, D. (2020). CHIMERA: combining ring-lwe-based fully homomorphic encryption schemes. *J. Math. Cryptol.*, 14(1):316–338.

Chen, H., Cammarota, R., Valencia, F., and Regazzoni, F. (2019). Plaidml-he: Acceleration of deep learning kernels to compute on encrypted data. In *37th IEEE International Conference on Computer Design, ICCD 2019, Abu Dhabi, United Arab Emirates, November 17-20, 2019*, pages 333–336. IEEE.

Cheon, J. H., Kim, A., Kim, M., and Song, Y. (2017). Homomorphic encryption for arithmetic of approximate numbers. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 409–437. Springer.

Chou, E., Beal, J., Levy, D., Yeung, S., Haque, A., and Fei-Fei, L. (2018). Faster cryptonets: Leveraging sparsity for real-world encrypted inference. *CoRR*, abs/1811.09953.

Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K., Maleki, S., Musuvathi, M., and Mytkowicz, T. (2019a). Chet: an optimizing compiler for fully-homomorphic neural-network inferencing. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 142–156.

Dathathri, R., Saarikivi, O., Chen, H., Laine, K., Lauter, K. E., Maleki, S., Musuvathi, M., and Mytkowicz, T. (2019b). CHET: an optimizing compiler for fully-homomorphic neural-network inferencing. In McKin-

- ley, K. S. and Fisher, K., editors, *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, pages 142–156. ACM.
- Gilad-Bachrach, R., Dowlin, N., Laine, K., Lauter, K. E., Naehrig, M., and Wernsing, J. (2016). Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In Balcan, M. and Weinberger, K. Q., editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 201–210. JMLR.org.
- Ioffe, S. and Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning*, pages 448–456. PMLR.
- Juvekar, C., Vaikuntanathan, V., and Chandrakasan, A. (2018). GAZELLE: A low latency framework for secure neural network inference. In Enck, W. and Felt, A. P., editors, *27th USENIX Security Symposium, USENIX Security 2018, Baltimore, MD, USA, August 15-17, 2018*, pages 1651–1669. USENIX Association.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krizhevsky, A. (2009). Learning multiple layers of features from tiny images. Technical report, CIFAR.
- LeCun, Y., Cortes, C., and Burges, C. J. MNIST handwritten digit database.
- Lu, W., Huang, Z., Hong, C., Ma, Y., and Qu, H. (2020). PEGASUS: bridging polynomial and non-polynomial evaluations in homomorphic encryption. *IACR Cryptol. ePrint Arch.*, 2020:1606.
- Mishra, P., Lehmkuhl, R., Srinivasan, A., Zheng, W., and Popa, R. A. (2020). Delphi: A cryptographic inference service for neural networks. *IACR Cryptol. ePrint Arch.*, 2020:50.
- Mohassel, P. and Rindal, P. (2018). Aby^3 : A mixed protocol framework for machine learning. In Lie, D., Mannan, M., Backes, M., and Wang, X., editors, *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security, CCS 2018, Toronto, ON, Canada, October 15-19, 2018*, pages 35–52. ACM.
- Regev, O. (2005). On lattices, learning with errors, random linear codes, and cryptography. In Gabow, H. N. and Fagin, R., editors, *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22-24, 2005*, pages 84–93. ACM.
- Riazi, M. S., Samragh, M., Chen, H., Laine, K., Lauter, K. E., and Koushanfar, F. (2019). XONN: xnor-based oblivious deep neural network inference. In Heninger, N. and Traynor, P., editors, *28th USENIX Security Symposium, USENIX Security 2019, Santa Clara, CA, USA, August 14-16, 2019*, pages 1501–1518. USENIX Association.
- SEAL (2020). Microsoft SEAL (release 3.6). <https://github.com/Microsoft/SEAL>. Microsoft Research, Redmond, WA.
- van Elsloo, T., Patrini, G., and Ivey-Law, H. (2019). Sealion: a framework for neural network inference on encrypted data. *CoRR*, abs/1904.12840.
- Viand, A., Jattke, P., and Hithnawi, A. (2021). Sok: Fully homomorphic encryption compilers. *arXiv preprint arXiv:2101.07078*.