

Théo Ryffel\*, Pierre Tholoniât, David Pointcheval, and Francis Bach

# AriaNN: Low-Interaction Privacy-Preserving Deep Learning via Function Secret Sharing

**Abstract:** We propose ARIANN, a low-interaction privacy-preserving framework for private neural network training and inference on sensitive data.

Our semi-honest 2-party computation protocol (with a trusted dealer) leverages function secret sharing, a recent lightweight cryptographic protocol that allows us to achieve an efficient online phase. We design optimized primitives for the building blocks of neural networks such as ReLU, MaxPool and BatchNorm. For instance, we perform private comparison for ReLU operations with a single message of the size of the input during the online phase, and with preprocessing keys close to  $4\times$  smaller than previous work. Last, we propose an extension to support  $n$ -party private federated learning.

We implement our framework as an extensible system on top of PyTorch that leverages CPU and GPU hardware acceleration for cryptographic and machine learning operations. We evaluate our end-to-end system for private inference between distant servers on standard neural networks such as AlexNet, VGG16 or ResNet18, and for private training on smaller networks like LeNet. We show that computation rather than communication is the main bottleneck and that using GPUs together with reduced key size is a promising solution to overcome this barrier.

**Keywords:** Multi Party Computation, Function Secret Sharing, Secure Comparison, Deep Learning

DOI 10.2478/popets-2022-0015

Received 2021-05-31; revised 2021-09-15; accepted 2021-09-16.

---

\*Corresponding Author: **Théo Ryffel:** INRIA, Département d’informatique de l’ENS, ENS, CNRS, PSL University, Paris, France, E-mail: theo.ryffel@ens.fr

**Pierre Tholoniât:** Columbia University, New York, USA

**David Pointcheval:** Département d’informatique de l’ENS, ENS, CNRS, PSL University, INRIA, Paris, France

**Francis Bach:** INRIA, Département d’informatique de l’ENS, ENS, CNRS, PSL University, Paris, France

## 1 Introduction

The massive improvements of cryptography techniques for secure computation over sensitive data [20, 22, 37] have spurred the development of the field of privacy-preserving machine learning [3, 56]. Privacy-preserving techniques have become practical for concrete use cases, thus encouraging public authorities to use them to protect citizens’ data especially in healthcare applications [24, 35, 49].

However, tools are lacking to provide end-to-end solutions for institutions that have little expertise in cryptography while facing critical data privacy challenges. A striking example is hospitals, which handle large amounts of data while having relatively constrained technical teams. Secure multi-party computation (SMPC) is a promising technique that can be efficiently integrated into machine learning workflows to ensure data and model privacy, while allowing multiple parties or institutions to participate in a joint project. In particular, SMPC provides intrinsic shared governance: because data is shared, none of the parties can decide alone to reconstruct it.

**Use case.** The main use case driving our work is the collaboration between a healthcare institution and an AI company. The healthcare institution, an hospital for example, acts as the data owner and the AI company as the model owner. The collaboration consists of either training the model with labelled data or using a pre-trained model to analyze unlabelled data. Training can possibly involve several data owners, as detailed in Section 5. Since the model can be a sensitive asset (in terms of intellectual property, strategic asset or regulatory and privacy issues), it cannot be trained directly on the data owner(s) machines using techniques like federated learning [11, 38]: it could be stolen or reverse-engineered [26, 31].

We will assume that the parties involved in the computation are located in different regions, and that they can communicate large amounts of information over the network with a reasonable latency (70ms for example). This corresponds to the *Wide Area Network (WAN)* setting, as opposed to the *Local Area Network (LAN)* setting where parties are typically located in the same

data center and communicate with low latency (typically  $<1\text{ms}$ ). Second, parties are *honest-but-curious*, [28, Chapter 7.2.2] and care about their reputation. Hence, they have little incentive to deviate from the original protocol, but they will use any information available in their own interest.

**Contributions.** By leveraging function secret sharing (FSS) [14, 15], we propose a low-interaction framework for private deep learning which drastically reduces communication to a single round for basic machine learning operations, and achieves the first private evaluation benchmark on ResNet18 using GPUs.

- We improve upon existing work of [15] on function secret sharing to design compact and ready-to-implement algorithms for tensor private comparison, which is a building block for neural networks and can be run with a single round of communication. In particular, we significantly reduce the key size from roughly  $n(4\lambda + n)$  to  $n(\lambda + 2n)$ , which is a crucial parameter as the computation time is linear in the key size.
- We show how function secret sharing can be used in machine learning and provide privacy-preserving implementations of classical layers, including ReLU, MaxPool and BatchNorm, to allow secure evaluation and training of arbitrary models on private data.
- Last, we provide a GPU implementation and a hardware-accelerated CPU implementation of our private comparison protocol<sup>1</sup>. As AriaNN is built over PyTorch for other tensor operations, it can run either completely on the GPU or on the CPU. We show its practicality both in LAN and WAN settings by running private inference on CIFAR-10 and Tiny Imagenet with models such as AlexNet [40], VGG16 [58] and ResNet18 [30], and private training on MNIST using models like LeNet.

**Related work.** Related work in privacy-preserving machine learning encompasses SMPC and fully homomorphic encryption (FHE) techniques.

FHE only needs a single round of interaction but does not support efficient non-linearities. For example, nGraph-HE [9] and its extensions [8] build on the SEAL library [55] and provide a framework for secure evaluation that greatly improves on the CryptoNet seminal work [27], but it resorts to polynomials (like the square) for activation functions.

SMPC frameworks usually provide faster implementations using lightweight cryptography. MiniONN [44], DeepSecure [52] and XONN [50] use optimized garbled circuits [63] that allow very few communication rounds, but they do not support training and alter the neural network structure to speed up execution. Other frameworks such as ShareMind [10], SecureML [46], SecureNN [59], QUOTIENT [2] or more recently FALCON [60] rely on additive secret sharing and allow secure model evaluation and training. They use simpler and more efficient primitives, but require a large number of rounds of communication, such as 11 in [59] or  $5 + \log_2(n)$  in [60] (typically 10 with  $n = 32$ ) for ReLU. ABY [23], Chameleon [51] and more recently ABY<sup>3</sup> [45], CryptFlow [41] and [21] mix garbled circuits, additive or binary secret sharing based on what is most efficient for the operations considered. However, conversion between those can be expensive and they do not support training except ABY<sup>3</sup>. There is a current line of work including BLAZE [47], Trident [18] and FLASH [17] which improves over ABY<sup>3</sup> to reduce communication overheads: BLAZE and Trident achieve for example 4 rounds of communication for ReLU.

Last, works like Gazelle [33] combine FHE and SMPC to make the most of both, but conversion can also be costly.

Works on trusted execution environments are left out of the scope of this article as they require access to dedicated and expensive hardware [32].

A concurrent work from Boyle et al. [13] was made public shortly after ours. Their approach also provides improvement over previous algorithms for private comparison using function secret sharing, and their implementation results in the same number of rounds than ours and similar key size (approximately  $n(\lambda + n)$ , where  $n$  is the number of bits to encode the value, it accounts for correctness and is typically set to 32, and  $\lambda$  is the security parameter and usually equals 128). However, [13] is not intended for machine learning: they only provide an implementation of ReLU, but not of MaxPool, BatchNorm, Argmax or other classic machine learning components. In addition, as they do not provide experimental benchmarks or an implementation of their private comparison, we are not able to compare it to ours in our private ML framework. They avoid the negligible error rate that we study in Section 3, which has no impact in the context of machine learning as we show.

<sup>1</sup> The code is available at [github.com/LaRiffle/AriaNN](https://github.com/LaRiffle/AriaNN).

## 2 Background

**Notations.** All values are encoded on  $n$  bits and live in  $\mathbb{Z}_{2^n}$ . The bit decomposition of any element  $x$  of  $\mathbb{Z}_{2^n}$  into a bit string of  $\{0, 1\}^n$  is a bijection between  $\mathbb{Z}_{2^n}$  and  $\{0, 1\}^n$ . Therefore, bit strings generated by a pseudo random generator  $G$  are implicitly mapped to  $\mathbb{Z}_{2^n}$ . In addition, we interpret the most significant bit as a sign bit to map them in  $[-2^{n-1}, 2^{n-1} - 1]$ , notably in Algorithms 1, 2, 3, 4, 5, where the modulo operation makes the conversion between  $n$  bit strings and signed integers explicit.

The notation  $\llbracket x \rrbracket$  denotes 2-party additive secret sharing of  $x$ , i.e.,  $\llbracket x \rrbracket = (\llbracket x \rrbracket_0, \llbracket x \rrbracket_1)$  where the shares  $\llbracket x \rrbracket_j$  are random in  $\mathbb{Z}_{2^n}$ , are held by distinct parties and verify  $x = \llbracket x \rrbracket_0 + \llbracket x \rrbracket_1 \pmod{2^n}$ . In return,  $x[i]$  refers to the  $i$ -th bit of  $x$ . The comparison operator  $\leq$  is taken over the natural embedding of  $\mathbb{Z}_{2^n}$  into  $\mathbb{Z}$ .

### 2.1 Function Secret Sharing

Unlike classical data secret sharing, where a shared input  $\llbracket x \rrbracket$  is applied on a public  $f$ , function secret sharing applies a public input  $x$  on a private shared function  $\llbracket f \rrbracket$ . Shares or *keys*  $(\llbracket f \rrbracket_0, \llbracket f \rrbracket_1)$  of a function  $f$  satisfy  $f(x) = \llbracket f \rrbracket_0(x) + \llbracket f \rrbracket_1(x) \pmod{2^n}$  and they can be provided by a semi-trusted dealer. Both approaches output a secret shared result.

Let us take an example: say Alice and Bob respectively have shares  $\llbracket y \rrbracket_0$  and  $\llbracket y \rrbracket_1$  of a private input  $y$ , and they want to compute  $\llbracket y \leq 0 \rrbracket$ . They first mask their shares using a random mask  $\llbracket \alpha \rrbracket$ , by computing  $\llbracket y \rrbracket_0 + \llbracket \alpha \rrbracket_0$  and  $\llbracket y \rrbracket_1 + \llbracket \alpha \rrbracket_1$ , and then reveal these values to reconstruct  $x = y + \alpha$ . Next, they apply this public  $x$  on their function shares  $\llbracket f_\alpha \rrbracket_j$  of  $f_\alpha : x \rightarrow (x \leq \alpha)$ , to obtain a shared output  $(\llbracket f_\alpha \rrbracket_0(x), \llbracket f_\alpha \rrbracket_1(x)) = \llbracket f_\alpha(y + \alpha) \rrbracket = \llbracket (y + \alpha) \leq \alpha \rrbracket = \llbracket y \leq 0 \rrbracket$ . [14, 15] have shown the existence of such function shares for comparison which perfectly hide  $y$  and the result. From now on, to be consistent with the existing literature, we will denote the function keys  $(k_0, k_1) := (\llbracket f \rrbracket_0, \llbracket f \rrbracket_1)$ .

Note that for a perfect comparison,  $y + \alpha$  should not wrap around and become negative. Because typically values of  $y$  used in practice in machine learning are small compared to the  $n$ -bit encoding amplitude with typically  $n = 32$ , the failure rate is less than one comparison in a million, as detailed in Section 3.2.

### 2.2 2-Party Computation in the Preprocessing Model

Preprocessing is performed during an offline phase by a trusted third party that builds and distributes the function keys to the 2 parties involved in future computation. This is standard in function secret sharing, and as mentioned by [16], in the absence of such trusted dealer, the keys can alternatively be generated via an interactive secure protocol that is executed offline, before the inputs are known. This setup can also be found in other privacy-preserving machine learning frameworks including SecureML [46]. This trusted dealer is not active during the online phase, and he is unaware of the computation the 2 parties intend to execute. In particular, as we are in the honest-but-curious model, it is assumed that no party colludes with the dealer. In practice, such third party would typically be an institution concerned about its reputation, and it could be easy to check that preprocessed material is correct using a *cut-and-choose* technique [65]. For example, the third party produces  $n$  keys for private comparison. The 2 parties willing to do the private computation randomly check some of them: they extract from their keys  $s_0, s_1$  and also reconstruct  $\alpha$  from  $\llbracket \alpha \rrbracket_j, j \in \{0, 1\}$ . They can then derive the computations of KeyGen and verify that the correlated randomness of the keys was correct. They can then use the remaining keys for the private computation.

### 2.3 Security Model of the Function Secret Sharing Protocol

We consider security against *honest-but-curious* adversaries, i.e., parties following the protocol but trying to infer as much information as possible about others' input or function share. This is a standard security model in many SMPC frameworks [7, 10, 51, 59] and is aligned with our main use case: parties that would not follow the protocol would face major backlash for their reputation if they got caught. The security of our protocols relies on indistinguishability of the function shares, which informally means that the shares received by each party are computationally indistinguishable from random strings. More formally, we introduce the following definitions from [15].

**Definition 2.1** (FSS: Syntax). A (2-party) *function secret sharing (FSS) scheme* is a pair of algorithms (KeyGen, Eval) with the following syntax:

- $\text{KeyGen}(1^\lambda, \hat{f})$  is a PPT *key generation* algorithm, which on input  $1^\lambda$  (security parameter) and  $\hat{f} \in \{0, 1\}^*$ , description of a function  $f : \mathbb{Z}_{2^n} \rightarrow \mathbb{Z}_{2^n}$ , outputs a pair of keys  $(k_0, k_1)$ .
- $\text{Eval}(i, k_i, x)$  is a polynomial-time *evaluation* algorithm, which on input  $i \in \{0, 1\}$  (party index),  $k_i$  (the  $i$ -th function key) and  $x \in \mathbb{Z}_{2^n}$ , outputs  $\llbracket f \rrbracket_i(x) \in \mathbb{Z}_{2^n}$  (the  $i$ -th share of  $f(x)$ ).

**Definition 2.2** (FSS: Correctness and Security). We say that  $(\text{KeyGen}, \text{Eval})$  as in Definition 2.1 is a *FSS scheme for a family of function  $\mathcal{F}$*  if it satisfies the following requirements:

- **Correctness:** For all  $f : \mathbb{Z}_{2^n} \rightarrow \mathbb{Z}_{2^n} \in \mathcal{F}$ ,  $\hat{f}$  a description of  $f$ , and  $x \in \mathbb{Z}_{2^n}$ , if  $(k_0, k_1) \leftarrow \text{KeyGen}(1^\lambda, \hat{f})$  then  $\Pr[\text{Eval}(0, k_0, x) + \text{Eval}(1, k_1, x) = f(x)] = 1$ .
- **Security:** For each  $i \in \{0, 1\}$ , there is a PPT algorithm  $\text{Sim}_i$  (simulator), such that for every infinite sequence  $(\hat{f}_j)_{j \in \mathbb{N}}$  of descriptions of functions from  $\mathcal{F}$  and polynomial size input sequence  $x_j$  for  $f_j$ , the outputs of the following experiments **Real** and **Ideal** are computationally indistinguishable:
  - **Real <sub>$j$</sub>**  :  $(k_0, k_1) \leftarrow \text{KeyGen}(1^\lambda, \hat{f}_j)$  ; Output  $k_i$
  - **Ideal <sub>$j$</sub>**  : Output  $\text{Sim}_i(1^\lambda)$

[15] has proved the existence of efficient FSS schemes in particular for equality. Such protocols and the ones that we derive from this work are proved to be secure against semi-honest adversaries, and as mentioned by [16], they could be extended to guarantee *security with abort* against malicious adversaries using MAC authentication [22], which means that the protocol would abort if parties deviated from it.

## 2.4 General Security Guarantees and Threats

The 2-party interaction for private prediction, i.e. when the model is already trained, is an example of *Encrypted Machine Learning as a Service (EMLaaS)*. In this scenario, as stated above, even a malicious model owner could not disclose information about the private inputs or predictions. However, it could use a different model where the weights have been modified to make poor or biased predictions. It is difficult for the data owner to realize that the model owner is misbehaving or using a model whose performance is inferior to what it claims, and this is an issue users also have with standard *Machine Learning as a Service (MLaaS)*. Proving

that the computation corresponds to a *certified* given model would require to commit the model and would be costly. On the other side, the information obtained by the data owner about the model (i.e. the prediction on a given input) is the same as in MLaaS. Model inversion techniques [64] can leverage multiple calls to the model to try to build a new model with similar performance. There are not many defenses against this, except limiting access to the model, which is usually the case in MLaaS where data owners are given a quota of requests. Also, attacks like membership inference [57] or reverse-engineering [26, 31] methods could be used to unveil information about the dataset on which the model was originally trained. Using *differential privacy* [1, 25] during the initial training of the model can provide some guarantees [48] against these threats, but it has a trade-off between privacy and utility as differentially private models usually have poorer performance.

Beyond evaluation, the case of fully-encrypted training can also expose the parties to some threats. The most common one is *data poisoning* and consists of the data owner undermining the training by providing irrelevant data or labels that are wrong or biased [5]. This attack however does not affect privacy. In return, if the model owner gets the final model in plaintext at the end of the training, the privacy of the data owner is at risk because the model owner could use the aforementioned techniques to get information about the training data. Using differential privacy during the private training is important to mitigate this privacy leakage, and should also be applied in a  $n$ -party training setting.

All these threats must be taken seriously when building production-ready systems. However, they are independent of the function secret sharing protocol and can be addressed separately by combining our work with differential privacy libraries for deep learning.

## 3 Function Secret Sharing Primitives


Our algorithms for private equality and comparison are built on top of the work of [15], so the security assumptions are the same as in this article. We first present an algorithm for equality which is very close to the one of [15] but which is used as a basis to build the comparison protocol. We then describe the private comparison protocol, which improves over the work of [15] on Distributed Interval Functions (DIF) by specializing on the operations needed for neural network evaluation or

training. In particular, we are able to reduce the function key size from roughly  $n(4\lambda + n)$  to  $n(\lambda + 2n)$ .

### 3.1 Equality Test

We start by describing private equality as introduced by [15], which is slightly simpler than comparison and gives useful hints about how comparison works. The equality test consists in comparing a public input  $x$  to a private value  $\alpha$ . Evaluating the input using the function keys can be viewed as walking a binary tree of depth  $n$ , where  $n$  is the number of bits of the input (typically 32). Among all the possible paths, the path from the root down to  $\alpha$  is called the *special path*. Figure 1 illustrates this tree and provides a compact representation which is used by our protocol, where we do not detail branches for which all leaves are 0. Evaluation goes as follows: two evaluators are each given a function key which includes a distinct initial random state  $(s, t) \in \{0, 1\}^\lambda \times \{0, 1\}$ . Each evaluator starts from the root, at each step  $i$  goes down one node in the tree and updates his state depending on the bit  $x[i]$  using a common *correction word*  $CW^{(i)} \in \{0, 1\}^{2(\lambda+1)}$  from the function key. At the end of the computation, each evaluator outputs  $t$ . As long as  $x[i] = \alpha[i]$ , the evaluators stay on the special path and because the input  $x$  is public and common to them, they both follow the same path. If a bit  $x[i] \neq \alpha[i]$  is met, they leave the special path and should output 0; else, they stay on it all the way down, which means that  $x = \alpha$  and they should output 1.

**Intuition.** The main idea is that while they are on the special path, evaluators should have states  $(s_0, t_0)$  and  $(s_1, t_1)$  respectively, such that  $s_0$  and  $s_1$  are i.i.d. and  $t_0 \oplus t_1 = 1$ . When they leave it, the correction word should act to have  $s_0 = s_1$  but still indistinguishable from random and  $t_0 = t_1$ , which ensures  $t_0 \oplus t_1 = 0$ . To reconstruct the result in plaintext, each evaluator should output its  $t_j$  and the result will be given by  $t_0 \oplus t_1$ . The formal description of the protocol is given below and is composed of two parts: first, in Algorithm 1, the KeyGen algorithm consists of a preprocessing step to generate the functions keys, and then, in Algorithm 2, Eval is run by two evaluators to perform the equality test. It takes as input the private share held by each evaluator and the function key that they have received. They use  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2(\lambda+1)}$ , a pseudorandom generator (PRG), where the output set is  $\{0, 1\}^{\lambda+1} \times \{0, 1\}^{\lambda+1}$ , and operations modulo  $2^n$  implicitly convert back and forth  $n$ -bit strings into integers.



**Fig. 1.** (Above) Binary decision tree with the special path for  $n = 3$ . Given an input  $x = x[1] \dots x[n]$ , at each level  $i$ , one should take the path labeled by the value in the square equal to the bit value  $x[i]$ . (Below) Flat representation of the tree.

**Correctness.** Intuitively, the correction words  $CW^{(i)}$  are built from the expected state of each evaluator on the special path, i.e., the state that each should have at each node  $i$  if it is on the special path given some initial state. During evaluation, a correction word is applied by an evaluator only when it has  $t = 1$ . Hence, on the special path, the correction is applied only by one evaluator at each bit. If at step  $i$ , the evaluator stays on the special path, the correction word compensates the current states of both evaluators by xor-ing them with themselves and re-introduces a pseudorandom value  $s$  (either  $s_0^R \oplus s_1^R$  or  $s_0^L \oplus s_1^L$ ), which means the xor of their states is now  $(s, 1)$  but those states are still indistinguishable from random.

On the other hand, if  $x[i] \neq \alpha[i]$ , the new state takes the other half of the correction word, so that the xor of the two evaluators states is  $(0, 0)$ . From there, they have the same states and both have either  $t = 0$  or  $t = 1$ . They will continue to apply the same corrections at each step and their states will remain the same, meaning that  $t_0 \oplus t_1 = 0$ . A final computation is performed to obtain a shared  $\llbracket T \rrbracket$  modulo  $2^n$  of the result bit  $t = t_0 \oplus t_1 \in \{0, 1\}$ .

**Security.** From the privacy point of view, when the seed  $s$  is random,  $G(s)$  is indistinguishable from random (this is a pseudorandom bit-string). Each half is used either in the  $cw$  or in the next state, but not both. Therefore, the correction words  $CW^{(i)}$  do not contain information about the expected states and for  $j = 0, 1$ , the output  $k_j$  is independently uniformly distributed with respect to

**Initialisation:** Sample random  $\alpha \xleftarrow{\$} \mathbb{Z}_{2^n}$   
 Sample random  $s_j^{(1)} \xleftarrow{\$} \{0, 1\}^\lambda$  and set  $t_j^{(1)} \leftarrow j$ , for  $j = 0, 1$

- 1 **for**  $i = 1..n$  **do**
- 2      $(s_j^L \parallel t_j^L, s_j^R \parallel t_j^R) \leftarrow G(s_j^{(i)}) \in \{0, 1\}^{\lambda+1} \times \{0, 1\}^{\lambda+1}$ , for  $j = 0, 1$
- 3     **if**  $\alpha[i]$  **then**  $cw^{(i)} \leftarrow (0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1)$
- 4     **else**  $cw^{(i)} \leftarrow (s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0)$ ;
- 5      $CW^{(i)} \leftarrow cw^{(i)} \oplus G(s_0^{(i)}) \oplus G(s_1^{(i)}) \in \{0, 1\}^{\lambda+1} \times \{0, 1\}^{\lambda+1}$
- 6      $state_j \leftarrow G(s_j^{(i)}) \oplus (t_j^{(i)} \cdot CW^{(i)}) = (state_{j,0}, state_{j,1})$ , for  $j = 0, 1$
- 7     Parse  $s_j^{(i+1)} \parallel t_j^{(i+1)} = state_{j,\alpha[i]} \in \{0, 1\}^{\lambda+1}$ , for  $j = 0, 1$
- 8  $CW^{(n+1)} \leftarrow (-1)^{t_1^{(n+1)}} \cdot (1 - s_0^{(n+1)} + s_1^{(n+1)}) \bmod 2^n$
- 9 **return**  $k_j \leftarrow \llbracket \alpha \rrbracket_j \parallel s_j^{(1)} \parallel CW^{(1)} \parallel \dots \parallel CW^{(n+1)}$ , for  $j = 0, 1$

**Algorithm 1:** KeyGen: function key generation for equality (from [15])

**Input:**  $(j, k_j, \llbracket y \rrbracket_j)$  where  $j \in \{0, 1\}$  refers to the evaluator id

- 1 Parse  $k_j$  as  $\llbracket \alpha \rrbracket_j \parallel s^{(1)} \parallel CW^{(1)} \parallel \dots \parallel CW^{(n+1)}$
- 2 Publish  $\llbracket \alpha \rrbracket_j + \llbracket y \rrbracket_j \bmod 2^n$  and get revealed  $x = \alpha + y \bmod 2^n$
- 3 Let  $t^{(1)} \leftarrow j$
- 4 **for**  $i = 1..n$  **do**
- 5      $state \leftarrow G(s^{(i)}) \oplus (t^{(i)} \cdot CW^{(i)}) = (state_0, state_1)$
- 6     Parse  $s^{(i+1)} \parallel t^{(i+1)} = state_{x[i]}$
- 7 **return**  $\llbracket T \rrbracket_j \leftarrow (-1)^j \cdot (t^{(n+1)} \cdot CW^{(n+1)} + s^{(n+1)}) \bmod 2^n$

**Algorithm 2:** Eval: evaluation of the function key for the equality test  $y = 0$  (from [15])

$\alpha$  and  $s_{1-j}^{(1)}$ , in a computational way. As a consequence, at the end of the evaluation, for  $j = 0, 1$ ,  $\llbracket T \rrbracket_j$  also follows a distribution independent of  $\alpha$ . Until the shared values are reconstructed, even a malicious adversary cannot learn anything about  $\alpha$  nor the inputs of the other player.

**Implementation.** Function keys should be computed by a third party dealer and sent to the evaluators in advance, which requires one extra communication of the size of the keys. We use the trick of [15] to reduce the size of each correction word in the keys, from  $2(1 + \lambda)$  to  $(2 + \lambda)$  by reusing the pseudo-random  $\lambda$ -bit string dedicated to the state used when leaving the special path for the state used for staying onto it, since for the latter state the only constraint is the pseudo-randomness of the bitstring. Regarding the PRG, we use a Matyas-Meyer-Oseas one-way compression function with an AES block cipher, as in [36] or [61]. We concatenate several fixed key block ciphers to achieve the desired output length:  $G(x) = E_{k_1}(x) \oplus x \parallel E_{k_2}(x) \oplus x \parallel \dots$ . Using AES helps us to benefit from hardware acceleration: we used the aesni Rust library for CPU execution and the csprng library of PyTorch for GPU. More de-


tails about implementation can be found in Appendix C.1.

## 3.2 Comparison

Our main contribution to the function secret sharing scheme is for the comparison function, which is intensively used in neural network to build non-polynomial activation functions like ReLU: we build on the idea of the equality test to provide a synthetic and efficient protocol whose structure is very close to the previous one, and improves upon the former DIF scheme of [15] by significantly reducing the key size.

### 3.2.1 Intuition

Instead of seeing the special path as a simple path, we can see it as a frontier for the zone in the tree where  $x \leq \alpha$ . To evaluate  $x \leq \alpha$ , we could evaluate all the paths on the left of the special path and then sum up the results, but this is highly inefficient as it requires



**Fig. 2.** (Above) Binary decision tree with all the paths corresponding to  $x \leq \alpha$  for  $n=3$ . (Below) Flat representation of the tree.

exponentially many evaluations. The key idea here is to evaluate all these paths at the same time, noting that each time one leaves the special path, it either falls on the left side (i.e.,  $x < \alpha$ ) or on the right side (i.e.,  $x > \alpha$ ). Hence, we only need to add an extra step at each node of the evaluation, where depending on the bit value  $x[i]$ , we output a leaf label which is 1 only if  $x[i] < \alpha[i]$  and all previous bits are identical. Only one label between the final label (which corresponds to  $x = \alpha$ ) and the leaf labels can be equal to one, because only a single path can be taken. Therefore, evaluators will return the sum of all the labels to get the final output.

### 3.2.2 Correctness

**Correctness of the comparison protocol.** Consider  $(k_0, k_1)$  generated by KeyGen (Algorithm 3) with a random offset  $\alpha \in \mathbb{Z}_{2^n}$ . Consider a public input  $x \in \mathbb{Z}_{2^n}$ . Let us show that  $\text{Eval}(0, k_0, x) + \text{Eval}(1, k_1, x) = (x \leq \alpha) \bmod 2^n$ , where  $(x \leq \alpha) \in \{0, 1\}$ . We add a subscript 0 or 1 to the variables of Algorithm 4 to identify the party to which they belong.

Consider  $i \in [1, n]$  such that the evaluators remained on the special path until  $i$  (i.e.  $\forall j < i, x[j] = \alpha[j]$ ). In particular,  $G(s_0^{(i)}) \oplus G(t_0^{(i)} \cdot CW^{(i)}) \oplus (s_1^{(i)}) \oplus (t_1^{(i)} \cdot CW^{(i)}) = cw^{(i)}$ . Let us study the 4 possible cases and show that 1)  $(out_{i,0} + out_{i,1} \bmod 2^n) \in \{0, 1\}$ ; 2)  $out_{i,0} + out_{i,1} = 1 \bmod 2^n \iff x[i] < \alpha[i]$ ; and 3) the evaluators stay on the special path if and only if  $x[i] = \alpha[i]$ .

- If  $x[i] = 0$ , we keep the left part of  $state'$  at line 4.
  - If  $\alpha[i] = 1$ , we have  $\tau_0^{(i+1)} \oplus \tau_1^{(i+1)} = 1$ . Thanks to line 10 of KeyGen, we have  $out_{i,0} + out_{i,1} = (\tau_0^{(i+1)} - \tau_1^{(i+1)}) CW_{leaf}^{(i)} + (\sigma_0^{(i+1)} - \sigma_1^{(i+1)}) = (1 - 2\tau_1^{(i+1)})(-1)^{\tau_1^{(i+1)}}(\sigma_1^{(i+1)} - \sigma_0^{(i+1)} + 1) + (\sigma_0^{(i+1)} - \sigma_1^{(i+1)}) = 1 \bmod 2^n$ . We also have  $t_0^{(i+1)} \oplus t_1^{(i+1)} = 0$  and  $s_0^{(i+1)} \oplus s_1^{(i+1)} = 0$ , so the evaluators leave the special path.
  - If  $\alpha[i] = 0$ , we use line 5 of KeyGen to generate  $cw^{(i)}$ , so  $\sigma_0^{(i+1)} \oplus \sigma_1^{(i+1)} = 0$  and  $\tau_0^{(i+1)} \oplus \tau_1^{(i+1)} = 0$ . Hence,  $out_{i,0} + out_{i,1} = (\tau_0^{(i+1)} - \tau_1^{(i+1)}) CW_{leaf}^{(i)} + (\sigma_0^{(i+1)} - \sigma_1^{(i+1)}) = 0 \bmod 2^n$ . We also have  $t_0^{(i+1)} \oplus t_1^{(i+1)} = 1$  and  $s_0^{(i+1)}, s_1^{(i+1)}$  stay on the special path.
- If  $x[i] = 1$ , we keep the right part of  $state'$  at line 4.
  - If  $\alpha[i] = 1$ , we use line 4 of KeyGen to generate  $cw^{(i)}$ , so  $\sigma_0^{(i+1)} \oplus \sigma_1^{(i+1)} = 0$  and  $\tau_0^{(i+1)} \oplus \tau_1^{(i+1)} = 0$ . Hence, similarly as the case where  $(x[i], \alpha[i]) = (0, 0)$ , we have  $out_{i,0} + out_{i,1} = 0 \bmod 2^n$  and the evaluators stay on the special path.
  - If  $\alpha[i] = 0$ , we use line 5 of KeyGen and get  $cw^{(i)} = ((s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, \sigma_0^L \oplus \sigma_1^L \parallel 1))$ . We keep the right part of  $state'$  at line 9 of KeyGen for  $CW_{leaf}^{(i)}$ . We use the same right part at line 5 of Eval, so we have  $\tau_0^{(i+1)} \oplus \tau_1^{(i+1)} = 1$ . Finally,  $out_{i,0} + out_{i,1} = (\tau_0^{(i+1)} - \tau_1^{(i+1)}) CW_{leaf}^{(i)} + (\sigma_0^{(i+1)} - \sigma_1^{(i+1)}) = (1 - 2\tau_1^{(i+1)})(-1)^{\tau_1^{(i+1)}}(\sigma_1^{(i+1)} - \sigma_0^{(i+1)} + 0) + (\sigma_0^{(i+1)} - \sigma_1^{(i+1)}) = 0 \bmod 2^n$ . We also have  $t_0^{(i+1)} \oplus t_1^{(i+1)} = 0$  and the evaluators leave the special path.

If the evaluators leave the special path at step  $i$ , their bistrings remain equal until the end of the evaluation:  $\forall j \in [i, n+1], s_0^{(j)} = s_1^{(j)}$  and  $\sigma_0^{(j)} = \sigma_1^{(j)}$ , so  $\forall j \in [i, n+1], out_{j,0} + out_{j,1} = 0 \bmod 2^n$ .

Finally, if the evaluators never leave the special path (i.e.  $x = \alpha$ ), we have  $\forall j \in [1, n], out_{j,0} + out_{j,1} = 0 \bmod 2^n$ , and  $out_{n+1,0} + out_{n+1,1} = 1$ . Indeed, step  $n+1$  is identical to the equality case (Algorithm 2).

In the end, the sum  $\llbracket T \rrbracket_j$  of the  $out_i$ 's is a share of 1 either if  $out_{n+1}$  was a share of 1 (i.e.  $x = \alpha$ ) or if one of the other  $out_i$  was a share of 1, which is possible only if  $\alpha[i] = 1$  and  $x[i] < \alpha[i]$  (i.e.  $x < \alpha$ ). Otherwise (i.e.  $x > \alpha$ ),  $\llbracket T \rrbracket_j$  is a share of 0.

**Failure rate of the sign protocol.** Algorithm 5 details how we build a sign protocol thanks to our com-

**Initialisation:** Sample random  $\alpha \xleftarrow{\$} \mathbb{Z}_{2^n}$   
Sample random  $s_j^{(1)} \xleftarrow{\$} \{0, 1\}^\lambda$  and set  $t_j^{(1)} \leftarrow j$ , for  $j = 0, 1$

- 1 **for**  $i = 1..n$  **do**
- 2     **for**  $j = 0, 1$  **do**
- 3          $((s_j^L \parallel t_j^L, s_j^R \parallel t_j^R), (\sigma_j^L \parallel \tau_j^L, \sigma_j^R \parallel \tau_j^R)) \leftarrow G(s_j^{(i)}) \in \{0, 1\}^{\lambda+1} \times \{0, 1\}^{\lambda+1} \times \{0, 1\}^{n+1} \times \{0, 1\}^{n+1}$
- 4         **if**  $\alpha[i]$  **then**  $cw^{(i)} \leftarrow ((0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1), (\sigma_0^R \oplus \sigma_1^R \parallel 1, 0^\lambda \parallel 0))$
- 5         **else**  $cw^{(i)} \leftarrow ((s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, \sigma_0^L \oplus \sigma_1^L \parallel 1));$
- 6          $CW^{(i)} \leftarrow cw^{(i)} \oplus G(s_0^{(i)}) \oplus G(s_1^{(i)})$
- 7         **for**  $j = 0, 1$  **do**
- 8              $state_j \leftarrow G(s_j^{(i)}) \oplus (t_j^{(i)} \cdot CW^{(i)}) = ((state_{j,0}, state_{j,1}), (state'_{j,0}, state'_{j,1}))$
- 9             Parse  $s_j^{(i+1)} \parallel t_j^{(i+1)} = state_{j,\alpha[i]}$  and  $\sigma_j^{(i+1)} \parallel \tau_j^{(i+1)} = state'_{j,1-\alpha[i]}$
- 10           $CW_{leaf}^{(i)} \leftarrow (-1)^{\tau_1^{(i+1)}} \cdot (\alpha[i] - \sigma_0^{(i+1)} + \sigma_1^{(i+1)}) \bmod 2^n$
- 11  $CW_{leaf}^{(n+1)} \leftarrow (-1)^{t_1^{(n+1)}} \cdot (1 - s_0^{(n+1)} + s_1^{(n+1)}) \bmod 2^n$
- 12 **return**  $k_j \leftarrow \llbracket \alpha \rrbracket_j \parallel s_j^{(1)} \parallel (CW^{(i)})_{i=1..n} \parallel (CW_{leaf}^{(i)})_{i=1..n+1}$ , for  $j = 0, 1$

**Algorithm 3:** KeyGen: function key generation for comparison  $x \leq \alpha$  (new)

**Input:**  $(j, k_j, x)$  where  $j \in \{0, 1\}$  refers to the evaluator id

- 1 Parse  $k_j$  as  $\llbracket \alpha \rrbracket_j \parallel s^{(1)} \parallel (CW^{(i)})_{i=1..n} \parallel (CW_{leaf}^{(i)})_{i=1..n+1}$
- 2 Let  $t^{(1)} \leftarrow j$
- 3 **for**  $i = 1..n$  **do**
- 4      $state \leftarrow G(s^{(i)}) \oplus (t^{(i)} \cdot CW^{(i)}) = ((state_0, state_1), (state'_0, state'_1))$
- 5     Parse  $s^{(i+1)} \parallel t^{(i+1)} = state_{x[i]}$  and  $\sigma^{(i+1)} \parallel \tau^{(i+1)} = state'_{x[i]}$
- 6      $out_i \leftarrow (-1)^j \cdot (\tau^{(i+1)} \cdot CW_{leaf}^{(i)} + \sigma^{(i+1)}) \bmod 2^n$
- 7  $out_{n+1} \leftarrow (-1)^j \cdot (t^{(n+1)} \cdot CW_{leaf}^{(n+1)} + s^{(n+1)}) \bmod 2^n$
- 8 **return**  $\llbracket T \rrbracket_j \leftarrow \sum_i out_i \bmod 2^n$

**Algorithm 4:** Eval: evaluation of the function key for comparison  $x \leq \alpha$  (new)

comparison primitive (Algorithm 4), following the secret sharing workflow introduced in Section 2.1. Our sign protocol can fail if  $y + \alpha$  wraps around and becomes negative. We cannot act on  $\alpha$  because it must be completely random to act as a perfect mask and to make sure the revealed  $x = y + \alpha \bmod 2^n$  does not leak any information about  $y$ , but the smaller  $y$  is, the lower the error probability will be. [16] suggests a method which uses 2 invocations of the protocol to guarantee perfect correctness but because it incurs an important runtime overhead, we rather show that the failure rate of our comparison protocol is very small and is reasonable in contexts that tolerate a few mistakes, as in machine learning. Consider  $y \in [-2^{n-1}, 2^{n-1} - 1]$ , a pair of comparison keys  $(k_0, k_1)$ , and note  $\widehat{\text{Sign}}(y) := \text{Sign}(0, k_0, \llbracket y \rrbracket_0) + \text{Sign}(1, k_1, \llbracket y \rrbracket_1) \bmod 2^n$  the reconstructed result of the sign protocol. We

have  $\Pr[\widehat{\text{Sign}}(y) \neq \mathbb{1}[y \leq 0]] = |y|/2^n \leq Y/2^n$  where  $Y$  is the maximum amplitude for  $|y|$ .

We quantify this failure rate on real world examples, namely on Network-2 and on the  $64 \times 64$  Tiny Imagenet version of VGG16, with a fixed precision of 3 decimals, and find respective failure rates of 1 in 4 millions comparisons and 1 in 100 millions comparisons, which is low compared to the number of comparisons needed for an evaluation, respectively  $\sim 10K$  and  $\sim 1M$ . In practice, such error rates do not affect the model accuracy, as Table 4 shows.

### 3.2.3 Security

The formal proof of security is provided in Appendix A.



**Input:**  $(j, k_j, \llbracket y \rrbracket_j)$  where  $j \in \{0, 1\}$  refers to the evaluator id

- 1 Parse the first  $n$  bits of  $k_j$  as  $\llbracket \alpha \rrbracket_j$
- 2 Publish  $\llbracket \alpha \rrbracket_j + \llbracket y \rrbracket_j$  and get revealed  $x = \alpha + y \bmod 2^n$
- 3 **return**  $\llbracket T \rrbracket_j \leftarrow \text{Eval}(j, k_j, x)$

**Algorithm 5:** Sign: protocol for  $\text{sign}(\llbracket y \rrbracket)$

### 3.2.4 Implementation and Communication Complexity

In all these computations modulo  $2^n$ , the bitstrings  $s_j^{(i)}$  and  $\sigma_j^{(i)}$  are respectively in  $\{0, 1\}^\lambda$  and  $\{0, 1\}^n$ , where we have typically  $\lambda = 128$  and  $n = 32$ . The PRG used here is  $G : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2(\lambda+1)+2(n+1)}$  where the output is seen as a pair of pairs of elements in  $(\{0, 1\}^{\lambda+1} \times \{0, 1\}^{\lambda+1}) \times (\{0, 1\}^{n+1} \times \{0, 1\}^{n+1})$ . For the right-hand part, we only need  $n$  bits instead of  $\lambda$  bits since the  $\sigma^{(i)}$  deriving from the PRG are not used for anything other than masking the  $n$ -bit output. This allows us to use fewer AES block ciphers in our PRG implementation and hence to achieve faster computation. In addition, because our comparison protocol works very similarly to the equality protocol, we can reuse the trick that consists of reusing randomness of the state corresponding of leaving the special area for the state corresponding of staying into it, as it does not compromise the fact that this state only needs to be pseudo-random. Thanks to this, we almost divide by 2 the size of the  $CW^{(i)}$  from  $2(\lambda+1)+2(n+1)$  to  $\lambda+2+n+2$ . Compared to the previous Distributed Interval Function (DIF) protocol of [15], our algorithm is not only much simpler as it does not require inspecting binary trees and searching for paths, but it also reduces significantly the key size from roughly  $n(4\lambda+n)$  to  $n(\lambda+2n+4)+\lambda+2n$  bits. This allows for faster transmission of keys over the network to the parties doing the evaluation.

## 4 Application to Deep Learning

We now apply these primitives to a private deep learning setup in which a model owner interacts with a data owner. The data and the model parameters are sensitive and are secret shared to be kept private. The shape of the input and the architecture of the model are however public, which is a standard assumption in secure deep learning [44, 46].

### 4.1 Additive Sharing Workflow with Preprocessing

All our operations are modular and follow this additive sharing workflow: inputs are provided secret shared and are masked with random values before being revealed. This disclosed value is then consumed with preprocessed function keys to produce a secret shared output. Each operation is independent of all surrounding operations, which is known as *circuit-independent preprocessing* [16] and implies that key generation can be fully outsourced without having to know the model architecture. This results in a fast runtime execution with a very efficient online communication, with a single round of communication and a message size equal to the input size for comparison.

Additionally, values need to be converted from float to fixed point precision before being secret shared. The fixed point representation allows one to store decimal values with some approximation using  $n$ -bits integers. For example, when using a fixed precision of 3, a decimal value  $x$  is stored as  $\lfloor x \cdot 10^3 \rfloor$  in  $\mathbb{Z}_{2^n}$ . Fixed precision is used to simplify operations like addition because the inputs can be summed up directly in  $\mathbb{Z}_{2^n}$ .

### 4.2 Common Machine Learning Operations

ReLU activation function is supported as a direct application of our comparison protocol, which we combine with a point wise multiplication. As mentioned in Section 2, this construction is not exact and is associated with an error rate which is below 1 in a million for typical ML computations. The comparison made in Table 4 between fixed point and private evaluation of pre-trained models shows that this error rate does not affect model accuracy.

**Matrix Multiplication (MatMul)**, as mentioned by [16], fits in this additive sharing workflow. We use Beaver triples [6] to compute  $\llbracket z \rrbracket = \llbracket x \cdot y \rrbracket$  from  $\llbracket x \rrbracket, \llbracket y \rrbracket$  and using a triple  $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket := \llbracket a \cdot b \rrbracket)$ , where all values are secret shared in  $\mathbb{Z}_{2^n}$ . The mask is here  $\llbracket (-a, -b) \rrbracket$

and is used to reveal  $(\delta, \epsilon) := (x - a, y - b)$ . The functional keys are the shares of  $(\llbracket a \rrbracket, \llbracket b \rrbracket, \llbracket c \rrbracket)$  and are used to compute  $\delta \cdot \llbracket b \rrbracket + \epsilon \cdot \llbracket a \rrbracket + \delta \cdot \epsilon + \llbracket c \rrbracket = \llbracket z \rrbracket$ . Matrix multiplication is identical but uses matrix Beaver triples [46].

**Convolution** can also be computed using Beaver triples. Using the previous notations, we can now consider  $y$  to be the convolution kernel, and the operation  $\cdot$  now stands for the convolution operator. We use this method for the CPU and GPU implementations, which enables us to use the PyTorch Conv2d function to compute the  $\cdot$  operation. Note that convolution can also be computed as a matrix multiplication using an unrolling technique as described in [19], but it incurs an overhead in terms of communication because the unrolled matrix is bigger than the original one when the stride is smaller than the kernel size. More details about unrolling can be found in Appendix C.2 with Figure 3.

**Argmax** is used to determine the predicted label for classification tasks (*i.e.* compute the index of the highest value of the last layer). Algorithm 6 shows how to compute this operator in a constant number of rounds using pairwise comparisons, in a fashion similar to [29]. This algorithm outputs the indices in the one-hot format, meaning that the output vector is of a similar shape to the input, and contains 1 where the maximum was found and 0 elsewhere. This protocol does not guarantee one-hot output: if the last layer outputs two identical maximum values, both will be retrieved. This sounds acceptable for machine learning evaluation as it informs that the model cannot choose between two classes. For training, the output signal only needs to be normalized. Probabilistic techniques are available to break ties, which only require an additional comparison.

In our algorithm, the first loop (line 2) requires  $m(m-1)$  parallel comparisons, and the second loop (line 4) requires  $m$  equality checks. Hence, the argmax uses 2 rounds of communication and sends  $O(m^2)$  values over the network. This is reasonable for a neural network where the number of outputs  $m$  is about 100 or less.

**MaxPool** can be implemented by combining the ideas of the unrolling-based convolution and the argmax: the matrix is first unrolled like in Figure 3 and the argmax of each row is then computed using parallel pairwise comparisons. This argmax is then multiplied with the row to get the maximum value, and the matrix is rolled back. These steps are illustrated in Figure 4 in Appendix C.3 and is formally described in Algorithm 7. It requires 3 rounds of communication, but we also provide an optimization when the kernel size  $k$  equals 2, which reduces

**Input:**  $\llbracket x_1 \rrbracket, \dots, \llbracket x_m \rrbracket$   
**Output:**  $\arg \max_{i \in [1, m]} x_i$

- 1 **for**  $j \in [1, m]$  **do**
- 2    $\llbracket s_j \rrbracket \leftarrow \sum_{i \neq j} \llbracket x_i - x_j \leq 0 \rrbracket$
- 3 **for**  $j \in [1, m]$  **do**
- 4    $\llbracket \delta_j \rrbracket \leftarrow \llbracket s_j = m - 1 \rrbracket$
- 5 **return**  $\llbracket \delta_1 \rrbracket, \dots, \llbracket \delta_m \rrbracket$

**Algorithm 6:** Argmax functionality using FSS

the computation complexity by a factor  $4 \times$  but uses an additional round of communication, and is very useful for some deep models such as VGG16.

**Input:**  $\llbracket \mathbf{X} \rrbracket = (\llbracket x_{i,j} \rrbracket)_{i,j=1 \dots m}$   
**Output:**  $\llbracket \text{MaxPool}(\mathbf{X}, k) \rrbracket$

- 1 Set  $n = \lfloor (m - k) / s + 1 \rfloor$
- 2 Define  $\llbracket \mathbf{X}^{\text{unrolled}} \rrbracket$  of shape  $n^2 \times k^2$
- 3 Define  $\llbracket \vec{y} \rrbracket$  of size  $n^2$
- 4 **for**  $i, j \in [0, s, 2s, \dots, m - k]$  **do**
- 5    $\llbracket x_{i,j}^{\text{unrolled}} \rrbracket = (\llbracket x_{i,j} \rrbracket, \dots, \llbracket x_{i+k,j+k} \rrbracket)$
- 6 **for**  $i \in [0, n^2 - 1]$  **do**
- 7    $\llbracket y_i \rrbracket \leftarrow \langle \llbracket \vec{x}_i \rrbracket, \text{Argmax}(\llbracket \vec{x}_i \rrbracket) \rangle$
- 8   where  $\llbracket \vec{x}_i \rrbracket = \llbracket x_{i,0} \rrbracket, \dots, \llbracket x_{i,k^2} \rrbracket$
- 9 **return**  $\vec{y}$  reshaped as a  $n \times n$  matrix.

**Algorithm 7:** MaxPool functionality using FSS.  $k$  is the kernel size, the stride is fixed to 2, padding to 0 and dilation to 1.

**BatchNorm** is implemented using Newton’s method as in [60] to implement the square inverse of the variance, as computing batch normalization exactly in a private way is very costly [59]. Given an input  $\vec{x} = (x_0, \dots, x_{m-1})$  with mean  $\mu$  and variance  $\sigma^2$ , we return  $\gamma \cdot \hat{\theta} \cdot (\vec{x} - \mu) + \beta$ . Variables  $\gamma$  and  $\beta$  are learnable parameters and  $\hat{\theta}$  is the estimate inverse of  $\sqrt{\sigma^2 + \epsilon}$  with  $\epsilon \ll 1$  and is computed iteratively as such:

$$\theta_{i+1} = \theta_i \cdot \frac{(C + 1) - (\sigma^2 + \epsilon) \cdot \theta_i^2}{C}$$

Compared to [60], we do not make any costly initial approximation, therefore instead of  $C = 2$  which corresponds to the classic Newton’s method, we use higher values of  $C$  (like  $C = 6$  for the intermediate layers) which can reduce the convergence speed of the method but spares the initialisation cost.

The requirements on the approximation depend whether we are doing training or evaluation. If we are

**Table 1.** Accuracy of training ResNet18 on the Hymenoptera classification task with exact or approximated BatchNorm (BN)

BatchNorm	init. with last batch	Newton iterations	Accuracy	Average comm. rounds per BN
Exact	-	-	93.59	-
Approx.	True	3	89.15	9
Approx.	False	20	88.24	60
Approx.	False	10	84.97	30
Approx.	False	3	60.13	9

evaluating a pre-trained secret-shared neural network, having a very precise approximation is crucial, especially if the model is deep like ResNet18. Indeed, the deeper the model is, the more errors in the BatchNorm layers will propagate in the model and make it unusable. However, if the model has a running mean and variance which is the default for PyTorch, we only need to compute once the square inverse of the running variance at the beginning of the computation.

For training however, we can use less precise approximations, since the goal of the batch normalization layer is to normalize the signal and this does not need to be done exactly as we show. We have found it very useful to reuse the result of the computation on the previous batch as an initial guess for the next batch. Moreover, we observe that for deep networks such as ResNet18, we can reduce the number of iterations of the Newton method from 4 to only 3 compared to [60], except of the first batch (which does not have a proper initialisation), and for the initial and last BatchNorm layers. For those layers, which either suffer from a too high or too low variance, we increase the number of iterations. For all layers, typical relative error never exceeds 5% and moderately affects learning capabilities, as our analysis on ResNet18 shows in Table 1. We train the model on the Hymenoptera binary classification task<sup>2</sup> using different approximated BatchNorm layers for which we report the associated number of rounds per layer when computed in a private way. More details about our experiments on ResNet18 can be found in Appendix C.5.

Table 2 summarizes the online communication cost of each ML operation presented above, and shows that basic operations such as comparison have a very efficient

online communication. We also report results from [60] which achieve good experimental performance.

### 4.3 Training Phase Using Autograd

These operations are sufficient to evaluate real world models in a fully private way. To also support private training of these models, we need to perform a private backward pass. As we overload operations such as convolutions or activation functions, we cannot use the built-in autograd functionality of PyTorch. Therefore, we have used the custom autograd functionality of the PySyft library [54], where it should be specified how to compute the derivatives of the operations that we have overloaded. Backpropagation also uses the same basic blocks than those used in the forward pass, including our private comparison protocol. Therefore, the training procedure `Train` described in Algorithm 8 closely follows the steps of plaintext training, except that the interactions between the secret shared data and model parameters use the protocols we have described in Section 4.2.

**Input:**  $\llbracket x \rrbracket, \llbracket y_{real} \rrbracket, \llbracket \theta \rrbracket$   
**Output:**  $\llbracket \hat{\theta} \rrbracket$

- 1 `opt = Optim( $\llbracket \theta \rrbracket$ )`
- 2  `$\llbracket y_{pred} \rrbracket = \text{Forward}(\llbracket \theta \rrbracket, \llbracket x \rrbracket)$`
- 3  `$\ell = \mathcal{L}(\llbracket y_{pred} \rrbracket, \llbracket y_{real} \rrbracket)$`
- 4  `$\llbracket \nabla \theta \rrbracket = \text{Backward}(\ell, \llbracket \theta \rrbracket)$`
- 5  `$\llbracket \theta \rrbracket = \text{opt}(\llbracket \nabla \theta \rrbracket, \llbracket \theta \rrbracket)$`
- 6 **return**  $\llbracket \theta \rrbracket$

**Algorithm 8:** Train procedure, that uses data  $x$  to update the model  $\theta$ . Lines 2 – 5 often run on batches extracted from  $\llbracket x \rrbracket$  and hence are iterated until  $\llbracket x \rrbracket$  has been completely used. `Optim` refers here to the optimizer that implements (stochastic) gradient descent. `Forward` and `Backward` are the forward and backward passes of the model  $\theta$ .  $\mathcal{L}$  is the loss function (mean square error or cross entropy).

## 5 Extension to Private Federated Learning

This 2-party protocol between a model owner and a data owner can be extended to an  $n$ -party federated learning

<sup>2</sup> [https://download.pytorch.org/tutorial/hymenoptera\\_data.zip](https://download.pytorch.org/tutorial/hymenoptera_data.zip)

**Table 2.** Theoretical online communication complexity of our protocols. Input sizes into brackets are those of the layers’ parameters, where  $k$  stands for the kernel size and  $s$  for the stride. Communication is given in number of values transmitted, and should be multiplied by their size (typically 4 bytes). Missing entries mean that data was not available.

Protocol	Input size	Rounds			Online Communication		
		Ours	FALCON [60]	ABY <sup>3</sup> [45]	Ours	FALCON [60]	ABY <sup>3</sup> [45]
Equality	$m$	1	-	2	$m$	-	$\sim \lambda m$
Comparison	$m$	1	7	2	$m$	$2m$	$\sim \lambda m$
MatMul	$m_1 \times m_2, m_2 \times m_3$	1	1	1	$m_1 m_2 + m_2 m_3$	$m_1 m_3$	$m_1 m_3$
Linear	$m_1 \times m_2, \{m_2 \times m_3\}$	1	1	-	$m_1 m_2 + m_2 m_3$	$m_1 m_3$	-
Convolution	$m \times m, \{k, s\}$	1	1	-	$((m - k)/s + 1)^2 k^2 + k^2$	$\sim m^2 k^2$	-
ReLU	$m$	2	10	-	$3m$	$4m$	-
Argmax	$m$	2	-	-	$m^2$	-	-
MaxPool	$m \times m, \{k, s\}$	3	$12(k^2 - 1)$	-	$((m - k)/s + 1)^2 (k^4 + 2)$	$\sim 5m^2$	-
BatchNorm	$m \times m$	9	335	-	$18m^2$	$\sim 56m^2$	-

protocol where several *clients* contribute their data to a model owned by an orchestrator *server*. We assume that the clients have the same set of features but have different samples in their data sets. This approach is sometimes called *Horizontal Federated Learning* and is used widely, like in secure aggregation [12]. The idea is that the server sends a version of the model to all clients, so that all clients start training the same model in parallel using their own data. With a frequency that varies depending on the settings, the server aggregates the models produced by each clients and sends back the aggregated version to be further trained by all clients. This way, clients *federate* their effort to train a global model, without sharing their data. Compared to secure aggregation [12], we are less concerned with parties dropping before the end of the protocol (we consider institutions rather than phones), and we do not reveal the updated model at each aggregation or at any stage, hence providing better privacy.

Algorithm 9 shows one possible implementation of fully private federated learning using 2-party function secret sharing. It prevents collusion between at most  $k$  out of  $n$  clients, the threat being that a client receiving the share of another client during aggregation phase could collude with the server to help reconstructing the model contributed by this client, and infer information about its private data. This aggregation requires extra communication rounds but this is in practice negligible compared to the training procedure `Train` initiated between a server and a client. Note that other aggregation mechanisms could be used, including using  $n$ -party MPC protocols or homomorphic encryption, but we proposed masking as this is quite in line with the concept of FSS where we mask the private input with  $\alpha$ .

## 6 Experiments

In order to simplify comparison with existing work, we follow a setup very close to the work of [60]. The reason why we compare our work to [60] is that it provides the most extensive experiments of private training and evaluation we are aware of. We are aware that [60] also provides honest-majority malicious security, but we only report their results in the honest but curious setting (where they obtain the best runtimes). We assess private inference of several networks on the datasets MNIST [43], CIFAR-10 [39],  $64 \times 64$  Tiny Imagenet [53, 62] and  $224 \times 224$  Hymenoptera which is a subset of Imagenet, and we also benchmark private training on MNIST. More details about the datasets used can be found in Appendix E.1. More precisely, we assess 6 networks: a 3 layers fully-connected network (Network-1), a small convolutional network with maxpool (Network-2), LeNet [42], AlexNet [40], VGG16 [58] and ResNet18 [30] which to the best of our knowledge has never been studied before in private deep learning. The description of these networks is available in Appendix E.2.

Our implementation provides a Python interface and is tightly coupled with PyTorch to provide both the ease of use and the expressiveness of this library. To use our protocols that only work in finite groups like  $\mathbb{Z}_{2^{32}}$ , we convert our input values and model parameters to fixed precision. To do so, we rely on the PySyft library [54] which extends common deep learning frameworks including PyTorch with a communication layer for federated learning and supports fixed precision. The experiments are run on Amazon EC2 using m5d.4xlarge machines for CPU benchmarks and g4dn.4xlarge for GPU, both with 16 cores and 64GB of CPU RAM, and we report our results both in the LAN and in the WAN

**Input:** Model on the server  $S$ , initialized with parameters  $\theta$   
**Output:** Model trained using the data from the clients  $(C_i)_{i=1..n}$

- 1 **Initialisation**  $S$  secret shares  $\theta$  with the clients.
- 2  $[[\theta]] \leftarrow_{\text{share}} \theta$
- 3 **for**  $i \in [1, n]$  **do**
- 4      $[[\theta_i]] \leftarrow_{\text{copy}} [[\theta]]$
- 5      $S$  stores  $[[\theta_i]]_0$  and sends  $[[\theta_i]]_1$  to  $C_i$
- 6 **Training**  $S$  runs in parallel  $n$  training procedures.
- 7 **for**  $i \in [1, n]$  **do**
- 8      $[[\hat{\theta}_i]] \leftarrow \text{Train}([x_i], [y_i], [[\theta_i]])$ , where  $[x_i]$  and  $[y_i]$  are the data and corresponding labels from  $C_i$
- 9 **Aggregation** All updated models are aggregated with a scheme secure against collusion between the server and  $k$  clients.
- 10  $S$  computes  $[[\hat{\theta}]]_0 := \sum_{i=1..n} [[\hat{\theta}_i]]_0$
- 11  $C^* \leftarrow^{\$} (C_i)_{i=1..n}$
- 12 **for**  $i \in [1, n]$  **do**
- 13      $C_i$  generates  $k$  seeds and sends them to  $C_{i+1 \bmod n}, \dots, C_{i+k \bmod n}$
- 14      $C_i$  receives  $k$  seeds from  $C_{i-k \bmod n}, \dots, C_{i-1 \bmod n}$
- 15      $C_i$  derives  $k$  random masks  $(m_j)_{j=1..k}$  from its own seeds
- 16      $C_i$  derives  $k$  random masks  $(\hat{m}_j)_{j=1..k}$  from the seeds received
- 17      $C_i$  builds a global mask  $\mu_i = \sum_{j=1..k} m_j - \hat{m}_j$
- 18      $C_i$  sends  $[[\theta_i]]_1 + \mu_i$  to  $C^*$
- 19  $C^*$  receives  $[[\hat{\theta}]]_1 := \sum_{i=1..n} [[\theta_i]]_1 + \sum_{i=1..n} \mu_i = \sum_{i=1..n} [[\theta_i]]_1$ .
- 20  $C^*$  broadcasts  $[[\hat{\theta}]]_1$  to all clients.
- 21 Iterate **Training** and **Aggregation** using  $[[\hat{\theta}]]$  until the training is complete.
- 22 **return**  $[[\hat{\theta}]]$

**Algorithm 9:** Federated Learning algorithm using 2-party Function Secret Sharing

setting. Latency is of 70ms for the WAN setting and is considered negligible in the LAN setting. Last, all values are encoded on 32 bits.

## 6.1 Inference Time and Communication

Comparison of experimental runtimes should be taken with caution, as different implementations and hardware may result in significant differences even for the same protocol. We report our online inference runtimes in Table 3 and show that they compare favourably with existing work including [44–46, 59, 60]. For example, our CPU implementation of Network-1 outperforms all other studied frameworks by at least a factor  $2\times$  in the LAN setting and even more in the WAN setting. For larger networks such as AlexNet and VGG16, we have an execution time which is slightly higher than [60]. One reason for this can be that we use a Python interface to serialize messages and communicate between parties, while [60] uses exclusively C code. However, we are more communication-efficient than [60] for models

starting from LeNet, with a typical gain of 7% to 30% on CIFAR-10. Regarding the high advantage we have on AlexNet and  $64\times 64$  Tiny Imagenet, this is explained by the fact that [60] uses a modified and more complex AlexNet while we use the one from PyTorch. Details about our networks architecture is given in Appendix E.3.

Results are given for a batched evaluation with a default batch size of 128 to amortize the communication cost, as in other works compared here. For larger networks, we reduce the batch size to have the preprocessing material (including the function keys) fitting into RAM, which reduces the benefit of amortization. The exact values chosen are available in Appendix D.

We have also added the results of our GPU implementation, which offers a clear speed-up compared to CPU with an execution which is between 10% and 100% faster. While this already shows the usefulness of using GPUs, one could expect a greater speed-up. One reason is that classic GPUs currently offer 16GB of RAM, which is a clear limitation for our work where we store keys in RAM. Storing the keys on the CPU would come

**Table 3.** Comparison of the inference time between secure frameworks on several popular neural network architectures. FALCON, SecureNN, and ABY<sup>3</sup> are 3-party protocols. XONN and Gazelle are 2-party protocols. All protocols are evaluated in the honest-but-curious setting. For each network we report in order the runtime for computing the preprocessing (if any) using CPUs, the runtime for the online phase in LAN using CPUs, the runtime for the online phase in LAN using GPUs, the runtime for the online phase in WAN using CPUs, and the communication needed during the online phase. Runtime is given in seconds and communication in MB. Missing entries mean that data was not available.

Framework	Dataset	Network-1					Network-2					LeNet				
		LAN Prep.	LAN CPU	LAN GPU	LAN CPU	WAN Comm.	LAN Prep.	LAN CPU	LAN GPU	LAN CPU	WAN Comm.	LAN Prep.	LAN CPU	LAN GPU	LAN CPU	WAN Comm.
AriaNN	MNIST	0.002	0.004	0.002	0.043	0.022	0.028	0.041	0.024	0.133	0.28	0.041	0.055	0.035	0.143	0.43
FALCON	MNIST	-	0.011	-	0.990	0.012	-	0.009	-	0.760	0.049	-	0.047	-	3.06	0.74
SecureNN	MNIST	-	0.043	-	2.43	2.1	-	0.130	-	3.93	8.86	-	-	-	-	-
XONN	MNIST	-	0.130	-	-	4.29	-	0.150	-	-	32.1	-	-	-	-	-
Gazelle	MNIST	0	0.030	-	-	0.5	0.481	0.330	-	-	22.5	-	-	-	-	-
ABY <sup>3</sup>	MNIST	0.005	0.003	-	-	0.5	-	-	-	-	-	-	-	-	-	-
CrypTFlow	MNIST	-	0.008	-	-	-	-	0.034	-	-	-	-	0.058	-	-	-

Framework	Dataset	AlexNet					VGG16					ResNet18				
		LAN Prep.	LAN CPU	LAN GPU	LAN CPU	WAN Comm.	LAN Prep.	LAN CPU	LAN GPU	LAN CPU	WAN Comm.	LAN Prep.	LAN CPU	LAN GPU	LAN CPU	WAN Comm.
AriaNN	CIFAR-10	0.091	0.15	0.078	0.34	0.95	0.94	1.75	1.55	1.99	12.59	-	-	-	-	-
FALCON	CIFAR-10	-	0.043	-	0.13	1.35	-	0.79	-	1.27	13.51	-	-	-	-	-
AriaNN	64×64 ImageNet	0.27	0.33	0.20	0.48	1.75	3.42	7.51	6.83	8.00	53.11	-	-	-	-	-
FALCON	64×64 ImageNet	-	1.81	-	2.43	19.21	-	3.15	-	4.67	52.56	-	-	-	-	-
AriaNN	224×224 Hymenoptera	-	-	-	-	-	-	-	-	-	-	10.02	19.88	13.90	24.07	148

**Table 4.** Accuracy of pre-trained neural network architectures, evaluated over several datasets in plaintext, fixed precision and privately using FSS. Time for private evaluation in the LAN setting is also reported.

Model	Dataset	LAN time (h)	Accuracy		
			Private	Fix prec.	Public
Network-1	MNIST	0.01	98.2	98.2	98.2
Network-2	MNIST	0.18	99.0	99.0	99.0
LeNet	MNIST	0.24	99.2	99.3	99.3
AlexNet	CIFAR-10	0.60	70.3	70.3	70.3
AlexNet	64×64 ImageNet	0.48	38.3	38.6	38.6
VGG16	CIFAR-10	5.19	87.4	87.4	87.4
VGG16	64×64 ImageNet	9.97	55.2	56.0	55.9
ResNet18	Hymenoptera	0.95	94.7	94.7	95.3

at a marginal cost of importing them on the GPU during the online phase but would allow to use bigger batches and hence better amortize the computation.

## 6.2 Test Accuracy

Thanks to the flexibility of our framework, we can train each of these networks in plaintext and need only one line of code to turn them into private networks where all parameters are secret shared, or to fixed precision

**Table 5.** Accuracy of neural network architectures trained over several datasets in plaintext, fixed precision and privately using FSS. Time for private training in the LAN setting is given in hours per epoch.

Model	Dataset	LAN time per epoch (h)	Accuracy			Epochs
			Private	Fix prec.	Public	
Network-1	MNIST	0.78	98.0	98.0	98.2	15
Network-2	MNIST	2.8	98.3	99.0	99.0	10
LeNet	MNIST	4.2	99.2	99.2	99.3	10

networks where all parameters are converted to fixed precision but computation is still in plaintext. Comparing the performance of private models with their fixed precision version helps us to understand if fixed precision by itself reduces the accuracy of the model, and gives an estimate of the loss that is related to using secret shared computation.

We compare the accuracy of several pre-trained networks in these 3 modes in Table 4 by running a private evaluation with FSS, a fixed precision using only PySyft and a public evaluation where the model is not modified. We observe that accuracy is well preserved overall and that converting to fixed precision has no impact on the accuracy of the model. We have a small reduction in accuracy for the two private models evaluated on 64×64 ImageNet but it remains close to the plaintext baseline.

This gap can be explained by the fact that PySyft uses a basic and approximate private truncation after multiplication where truncation is directly applied on the shares, and by the error rate of our FSS comparison protocol. The drop in accuracy on ResNet18 is also minor and corresponds to a single mislabeled item.

If we degrade the encoding precision which by default considers values in  $\mathbb{Z}_{2^{32}}$ , or the fixed precision which is by default of 4 decimals, performance degrades as shown in Appendix B.

### 6.3 Training Accuracy

We have also assessed the ability of training neural networks from scratch in a private way using AriaNN. Private training is an end-to-end private procedure, which means the model or the gradients are never accessible in plaintext. We use stochastic gradient descent (SGD) with momentum, a simple but popular optimizer, and support several losses such as mean square error (used for Network-1) and cross entropy (used for Network-2 and LeNet). We report the runtime and accuracy obtained by training from scratch and evaluating several networks in Table 5, in plaintext, in fixed precision and in a fully private way, just as we did for inference. Note that because of the training setting, accuracy might not match best known results, but the training procedure is the same for all training modes which allows for fair comparison.

We observe that the training is done almost perfectly both in fixed precision and private mode compared to the plaintext counterpart. The only noticeable difference we observe is for Network-2, where the privately-trained model achieves 98.3% while 99.0% is expected. The fixed-precision accuracy which is 99.0% suggests that our autograd functionality is working properly, so the difference must be explained by the small failure rate of FSS. Training profiles show that the accuracy starts decreasing roughly after 3 epochs, while it is supposed to keep increasing smoothly up to the 10<sup>th</sup> epoch. Instability caused by some FSS failures could account for this behaviour. However, training on LeNet did not suffer from the same phenomenon.

Recently, [21] also reports accuracy results when training securely Network-1 on MNIST, using a 3-party semi-honest protocol that mixes [45] and [4]. They achieve 97.8 % of accuracy in 15 epochs with a runtime of only 33.8s per epoch in the LAN setting. However, they do not provide a detailed comparison between the accuracy achieved with private training, and clear-

text training. One major difference with our work is that we are more communication efficient. We only require 10.3MB of communication during the online phase while they use 33.8MB per epoch. In addition, they rely on ABY<sup>3</sup>, which mean they use much more interaction rounds, which could be costly in the WAN setting although this is not monitored by this work.

Training cannot complete in reasonable time for larger networks such as VGG16, which in practice might be fine-tuned rather than trained from scratch. Note that training time includes the time spent building the preprocessing material, as it is too large to be fully processed and stored in RAM in advance.

### 6.4 Computation and Communication Analysis

We have provided in Table 6 a small analysis of how the compute time can be decomposed. We use AlexNet on the Tiny Imagenet dataset as it is the biggest network on which we could use a batch size higher than 64 both on CPU and GPU and hence amortize the serialization and communication cost.

**Table 6.** Distribution of the compute time during inference of AlexNet on the Tiny Imagenet dataset using either CPUs or GPUs.

Processor	FSS	MatMul and Convolution	Serialization and Deser.	Other
CPU	16%	72%	4%	8%
	53ms	238ms	13ms	26ms
GPU	51%	39%	8%	2%
	102ms	78ms	16ms	4ms

Thanks to the efficiency of our Rust implementation, function secret sharing only accounts for 16% for the on-line runtime when we use CPUs, and most of the time is spent doing matrix multiplications and convolutions. This last part uses the underlying PyTorch functions on integers which are significantly slower than when they run on floats. This motivates us to use GPUs for which such operations are far more efficient. In the GPU setting, the distribution of time is indeed much more balanced, and having function secret sharing directly running on GPUs avoids going back and forth between the CPU and the GPU.

Regarding the trade-off between computation and communication time, we show in Table 7 that in the

WAN setting and using CPUs, computation appears to be the main bottleneck especially for bigger models. This also encourages us to further improve the GPU implementation, as any optimization of the computation efficiency will have an important impact on the overall runtime.

**Table 7.** Proportion of the overall runtime spent on computation versus communication, in the WAN setting using CPUs

Model	Dataset	Computation (%)	Comm. (%)
Network-1	MNIST	9	91
Network-2	MNIST	31	69
LeNet	MNIST	38	62
AlexNet	CIFAR-10	44	56
AlexNet	64×64 ImageNet	69	31
VGG16	CIFAR-10	88	12
VGG16	64×64 ImageNet	93	7
ResNet18	Hymenoptera	83	17

## 6.5 Discussion

Regarding experiments on larger networks, we could not use batches of size 128. This is mainly due to the size of the comparison function keys, which is currently proportional to the size of the input tensor, with a multiplicative factor of  $n\lambda$  where  $n = 32$  and  $\lambda = 128$ . Optimizing the function secret sharing protocol to reduce the size of those keys would allow to better amortize batched computations and would also reduce the runtime as we would manipulate smaller arrays during the private comparison. An interesting other improvement would be to run experiments on  $n = 16$  bits instead of 32. Classic ML frameworks like PyTorch or TensorFlow now support 16 bits encoding both on CPU and GPU.

We have proposed a first implementation of FSS on GPU, which can still be improved to reduce the memory footprint of they keys. Further efforts could be made to decrease it roughly by 50% to match the theoretical key size. In addition, and as the small difference between the LAN and the WAN runtime shows, especially for bigger networks, most of the time is now spent on computation. Therefore, optimizing computation on GPUs will have a direct impact on the overall efficiency of the inference or the training.

We have shown the relevance of using FSS for private training and evaluation of models in machine learning. Compared to concurrent works like [13], we have

shown that we have very competitive protocols, and that the failure rate of the comparison protocol has no impact for machine learning applications. Our protocol has been used in one recent work [34] where it was applied to the field of medical imaging on chest X-rays.

## 7 Conclusion

In this work, we improve over the best known protocols for private comparison using function secret sharing by reducing the keys size by almost a factor  $\times 4$ . We show how this new algorithm helps us implement efficient machine learning components and we provide constructions for ReLU and MaxPool with only 2 and 3 rounds of communication. Additionally, we show that AriaNN can implement a large diversity of neural networks, from convolutional networks to ResNet18, which are very competitive in terms of runtime and communication compared to existing work. Last, we provide an implementation of AriaNN which can run both on CPU and GPU, providing promising runtime improvements for the next generation of hardware accelerated privacy-preserving machine learning models.

## 8 Acknowledgments

We would like to thank Geoffroy Couteau, Chloé Héban and Loïc Estève for helpful discussions throughout this project. We are also grateful for the long-standing support of the OpenMined community and in particular its dedicated cryptography team, including George Muraru, Rasswanth S, Hrishikesh Kamath, Arturo Marquez, Yugandhar Tripathi, S P Sharan, Muhammed Abogazia, Alan Aboudib, Ayoub Benaissa, Sukhad Joshi and many others.

This work was supported in part by the French project FUI ANBLIC. The computing power was graciously provided by the French company ARKHN.

## References

- [1] Martin Abadi, Andy Chu, Ian Goodfellow, H Brendan McMahan, Ilya Mironov, Kunal Talwar, and Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 308–318, 2016.



- [2] Nitin Agrawal, Ali Shahin Shamsabadi, Matt J Kusner, and Adrià Gascón. Quotient: two-party secure neural network training and prediction. In *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*, pages 1231–1247, 2019.
- [3] Mohammad Al-Rubaie and J. Morris Chang. Privacy-preserving machine learning: Threats and solutions. *IEEE Security & Privacy*, 17(2):49–58, 2019.
- [4] Toshinori Araki, Jun Furukawa, Yehuda Lindell, Ariel Nof, and Kazuma Ohara. High-throughput semi-honest secure three-party computation with an honest majority. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 805–817, 2016.
- [5] Eugene Bagdasaryan, Andreas Veit, Yiqing Hua, Deborah Estrin, and Vitaly Shmatikov. How to backdoor federated learning. In *International Conference on Artificial Intelligence and Statistics*, pages 2938–2948. PMLR, 2020.
- [6] Donald Beaver. Efficient multiparty protocols using circuit randomization. In *Annual International Cryptology Conference*, pages 420–432. Springer, 1991.
- [7] Aner Ben-Efraim, Yehuda Lindell, and Eran Omri. Optimizing semi-honest secure multiparty computation for the internet. In *Proceedings of the Conference on Computer and Communications Security*, pages 578–590, 2016.
- [8] Fabian Boemer, Anamaria Costache, Rosario Cammarota, and Casimir Wierzynski. nGraph-HE2: A high-throughput framework for neural network inference on encrypted data. In *Proceedings of the 7th ACM Workshop on Encrypted Computing & Applied Homomorphic Cryptography*, pages 45–56, 2019.
- [9] Fabian Boemer, Yixing Lao, Rosario Cammarota, and Casimir Wierzynski. nGraph-HE: a graph compiler for deep learning on homomorphically encrypted data. In *Proceedings of the ACM International Conference on Computing Frontiers*, pages 3–13, 2019.
- [10] Dan Bogdanov, Sven Laur, and Jan Willemson. Sharemind: A framework for fast privacy-preserving computations. In *European Symposium on Research in Computer Security*, pages 192–206. Springer, 2008.
- [11] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kidon, Jakub Konecny, Stefano Mazzocchi, and H. Brendan McMahan. Towards federated learning at scale: System design. *arXiv preprint arXiv:1902.01046*, 2019.
- [12] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the Conference on Computer and Communications Security*, pages 1175–1191, 2017.
- [13] Elette Boyle, Nishanth Chandran, Niv Gilboa, Divya Gupta, Yuval Ishai, Nishant Kumar, and Mayank Rathee. Function secret sharing for mixed-mode and fixed-point secure computation. *Cryptology ePrint Archive: Report 2020/1392*, 2020.
- [14] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 337–367. Springer, 2015.
- [15] Elette Boyle, Niv Gilboa, and Yuval Ishai. Function secret sharing: Improvements and extensions. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 1292–1303, 2016.
- [16] Elette Boyle, Niv Gilboa, and Yuval Ishai. Secure computation with preprocessing via function secret sharing. In *Theory of Cryptography Conference*, pages 341–371. Springer, 2019.
- [17] Megha Byali, Harsh Chaudhari, Arpita Patra, and Ajith Suresh. Flash: fast and robust framework for privacy-preserving machine learning. *Proceedings on Privacy Enhancing Technologies*, 2020(2):459–480, 2020.
- [18] Harsh Chaudhari, Rahul Rachuri, and Ajith Suresh. Trident: Efficient 4pc framework for privacy preserving machine learning. In *27th Annual Network and Distributed System Security Symposium, NDSS*, pages 23–26, 2020.
- [19] Kumar Chellapilla, Sidd Puri, and Patrice Simard. High performance convolutional neural networks for document processing. In *International Workshop on Frontiers in Handwriting Recognition*, 2006.
- [20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachene. Faster fully homomorphic encryption: Bootstrapping in less than 0.1 seconds. In *international Conference on the Theory and Application of Cryptology and Information Security*, pages 3–33. Springer, 2016.
- [21] Anders Dalskov, Daniel Escudero, and Marcel Keller. Fantastic four: Honest-majority four-party secure computation with malicious security. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, August 2021.
- [22] Ivan Damgård, Valerio Pastro, Nigel Smart, and Sarah Zakarias. Multiparty computation from somewhat homomorphic encryption. In *Annual Cryptology Conference*, pages 643–662. Springer, 2012.
- [23] Daniel Demmler, Thomas Schneider, and Michael Zohner. Aby-a framework for efficient mixed-protocol secure two-party computation. In *NDSS*, 2015.
- [24] Tamara Dugan and Xukai Zou. A survey of secure multiparty computation protocols for privacy preserving genetic tests. In *International Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE)*, pages 173–182. IEEE, 2016.
- [25] Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.
- [26] Matt Fredrikson, Somesh Jha, and Thomas Ristenpart. Model inversion attacks that exploit confidence information and basic countermeasures. In *Proceedings of the Conference on Computer and Communications Security*, pages 1322–1333, 2015.
- [27] Ran Gilad-Bachrach, Nathan Dowlin, Kim Laine, Kristin Lauter, Michael Naehrig, and John Wernsing. Cryptonets: Applying neural networks to encrypted data with high throughput and accuracy. In *International Conference on Machine Learning*, pages 201–210, 2016.
- [28] Oded Goldreich. *Foundations of Cryptography: volume 2, Basic Applications*. Cambridge University Press, 2009.
- [29] Awni Hannun, Brian Knott, Shubho Sengupta, and Laurens van der Maaten. Privacy-preserving contextual bandits. *arXiv preprint arXiv:1910.05299*, 2019.
- [30] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceed-*

- ings of the Conference on Computer Vision and Pattern Recognition, pages 770–778, 2016.
- [31] Briland Hitaj, Giuseppe Ateniese, and Fernando Perez-Cruz. Deep models under the gan: information leakage from collaborative deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 603–618, 2017.
- [32] Tyler Hunt, Congzheng Song, Reza Shokri, Vitaly Shmatikov, and Emmett Witchel. Chiron: Privacy-preserving machine learning as a service. *arXiv preprint arXiv:1803.05961*, 2018.
- [33] Chiraag Juvekar, Vinod Vaikuntanathan, and Anantha Chandrakasan. {GAZELLE}: A low latency framework for secure neural network inference. In *USENIX Security Symposium 18*, pages 1651–1669, 2018.
- [34] Georgios Kaissis, Alexander Ziller, Jonathan Passerat-Palmbach, Théo Ryffel, Dmitrii Usynin, Andrew Trask, Ionésio Lima, Jason Mancuso, Friederike Jungmann, Marc-Matthias Steinborn, et al. End-to-end privacy preserving deep learning on multi-institutional medical imaging. *Nature Machine Intelligence*, pages 1–12, 2021.
- [35] Harmanjeet Kaur, Neeraj Kumar, and Shalini Batra. An efficient multi-party scheme for privacy preserving collaborative filtering for healthcare recommender system. *Future Generation Computer Systems*, 86:297–307, 2018.
- [36] Marcel Keller, Emmanuela Orsini, and Peter Scholl. Mascot: faster malicious arithmetic secure computation with oblivious transfer. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 830–842, 2016.
- [37] Marcel Keller, Valerio Pastro, and Dragos Rotaru. Overdrive: making spdz great again. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 158–189. Springer, 2018.
- [38] Jakub Konečný, H Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- [39] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The CIFAR-10 dataset. *online: <http://www.cs.toronto.edu/kriz/cifar.html>*, 55, 2014.
- [40] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 1097–1105, 2012.
- [41] Nishant Kumar, Mayank Rathee, Nishanth Chandran, Divya Gupta, Aseem Rastogi, and Rahul Sharma. Cryptflow: Secure tensorflow inference. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 336–353. IEEE, 2020.
- [42] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [43] Yann LeCun, Corinna Cortes, and C. J. Burges. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [44] Jian Liu, Mika Juuti, Yao Lu, and Nadarajah Asokan. Oblivious neural network predictions via minionn transformations. In *Proceedings of the Conference on Computer and Communications Security*, pages 619–631, 2017.
- [45] Payman Mohassel and Peter Rindal. Aby3: A mixed protocol framework for machine learning. In *Proceedings of the Conference on Computer and Communications Security*, pages 35–52, 2018.
- [46] Payman Mohassel and Yupeng Zhang. Secureml: A system for scalable privacy-preserving machine learning. In *Symposium on Security and Privacy (SP)*, pages 19–38. IEEE, 2017.
- [47] Arpita Patra and Ajith Suresh. Blaze: Blazing fast privacy-preserving machine learning. *arXiv preprint arXiv:2005.09042*, 2020.
- [48] Md Atiqur Rahman, Tanzila Rahman, Robert Laganière, Noman Mohammed, and Yang Wang. Membership inference attack against differentially private deep learning model. *Trans. Data Priv.*, 11(1):61–79, 2018.
- [49] Leonie Reichert, Samuel Brack, and Björn Scheuermann. Privacy-preserving contact tracing of covid-19 patients. *Cryptology ePrint*, (2020/375), 2020.
- [50] M Sadegh Riazi, Mohammad Samragh, Hao Chen, Kim Laine, Kristin Lauter, and Farinaz Koushanfar. {XONN}: Xnor-based oblivious deep neural network inference. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pages 1501–1518, 2019.
- [51] M Sadegh Riazi, Christian Weinert, Oleksandr Tkachenko, Ebrahim M Songhori, Thomas Schneider, and Farinaz Koushanfar. Chameleon: A hybrid secure computation framework for machine learning applications. In *Proceedings of the Asia Conference on Computer and Communications Security*, pages 707–721, 2018.
- [52] Bitu Darvish Rouhani, M. Sadegh Riazi, and Farinaz Koushanfar. Deepsecure: Scalable provably-secure deep learning. In *Proceedings of the 55th Annual Design Automation Conference*, pages 1–6, 2018.
- [53] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, and Michael Bernstein. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3):211–252, 2015.
- [54] Théo Ryffel, Andrew Trask, Morten Dahl, Bobby Wagner, Jason Mancuso, Daniel Rueckert, and Jonathan Passerat-Palmbach. A generic framework for privacy preserving deep learning. *arXiv preprint arXiv:1811.04017*, 2018.
- [55] Microsoft SEAL (release 3.0). <http://sealcrypto.org>, October 2018. Microsoft Research, Redmond, WA.
- [56] Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In *Proceedings of the Conference on Computer and Communications Security*, pages 1310–1321, 2015.
- [57] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18. IEEE, 2017.
- [58] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [59] Sameer Wagh, Divya Gupta, and Nishanth Chandran. Securenn: Efficient and private neural network training. *IACR Cryptology ePrint Archive*, 2018:442, 2018.
- [60] Sameer Wagh, Shruti Tople, Fabrice Benhamouda, Eyal Kushilevitz, Prateek Mittal, and Tal Rabin. Falcon: Honest-majority maliciously secure framework for private deep learn-

ing. *arXiv preprint arXiv:2004.02229*, 2020.

- [61] Frank Wang, Catherine Yun, Shafi Goldwasser, Vinod Vaikuntanathan, and Matei Zaharia. Splinter: Practical private queries on public data. In *14th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 17)*, pages 299–313, 2017.
- [62] Jiayu Wu, Qixiang Zhang, and Guoxi Xu. Tiny imagenet challenge. Technical report, Available: <http://cs231n.stanford.edu/reports/2017/pdfs/930.pdf>, 2017.
- [63] Andrew Chi-Chih Yao. How to generate and exchange secrets. In *Annual Symposium on Foundations of Computer Science*, pages 162–167, 1986.
- [64] Yuheng Zhang, Ruoxi Jia, Hengzhi Pei, Wenxiao Wang, Bo Li, and Dawn Song. The secret revealer: generative model-inversion attacks against deep neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 253–261, 2020.
- [65] Ruiyu Zhu, Yan Huang, Jonathan Katz, and Abhi Shelat. The cut-and-choose game and its application to cryptographic protocols. In *USENIX Security Symposium Security 16*, pages 1085–1100, 2016.

## A FSS Comparison Protocol - Security Proof

For this proof, we follow the same process than [15].

We prove that each party’s key  $k_j$  is pseudorandom. This is done via a sequence of hybrid distributions, where in each step we replace two correction words  $CW^{(i)}$  and  $CW_{leaf}^{(i)}$  within the key from being honestly generated to being random. In the initial game, all the correction words are as in the real distribution, and in the last game, they are all random. As every gaps are indistinguishable for any polynomially-bounded adversary, the real distribution is indistinguishable from random: this proves the pseudo-randomness of the keys.

The high-level argument for security will go as follows. Each party  $j \in \{0, 1\}$  begins with a share  $[\alpha]_j$  and a random seed  $s_j^{(1)}$  that are completely unknown to the other party. In each level of key generation (for  $i = 1$  to  $n$ ), the parties apply a PRG to their seed  $s_j^{(i)}$  to generate 8 items: namely, 2 seeds  $s_j^L, s_j^R$ , 2 bits  $t_j^L, t_j^R$ , 2  $n$ -bits values  $\sigma_j^L, \sigma_j^R$  and 2 other bits  $\tau_j^L, \tau_j^R$ . This process is always done on a seed which appears completely random given the view of the other party. Hence, the security of the PRG guarantees that the 8 resulting values appear similarly random given the view of the other party. The  $i$ th level correction word  $CW^{(i)}$  will “use up” the secret randomness of 3 of the 4 first pieces: the two bits  $t_j^L, t_j^R$ , and the seed  $s_j^{Lose}$  corresponding to the direction exiting the special path i.e.  $Lose = L$  if  $\alpha[i] = 1$  and  $Lose = R$  if  $\alpha[i] = 0$ . However, given this  $CW^{(i)}$ , the remaining seed  $s_j^{Keep}$  for  $Keep \neq Lose$  is still unpredictable to the other party, as it is kept hidden. Similarly, the  $i$ th level correction word  $CW_{leaf}^{(i)}$  uses up the secret randomness of the 4 last pieces,  $\sigma_j^L, \sigma_j^R$  and  $\tau_j^L, \tau_j^R$ , and appears random given the view of the other party. The argument is then continued in similar fashion to the next level, which uses  $s_j^{Keep}$  as an input to the PRG.

For each  $i \in \{0, 1, \dots, n + 1\}$ , we will consider a hybrid distribution  $Hyb_i$  defined roughly as follows, for  $j \in \{0, 1\}$ :

1.  $s_j^{(1)} \xleftarrow{\$} \{0, 1\}^\lambda$  chosen at random (honestly), and  $t_j^{(1)} = j$ .
2.  $CW^{(1)}, \dots, CW^{(i)} \leftarrow \{0, 1\}^{2(\lambda+n+2)}$  and  $CW_{leaf}^{(1)}, \dots, CW_{leaf}^{(i)} \leftarrow \{0, 1\}^n$  chosen at random.
3. For  $k < i$ ,  $s_j^{(k+1)} || t_j^{(k+1)}, \sigma_j^{(k+1)} || \tau_j^{(k+1)}$  computed honestly, as a function of  $s_j^{(0)} || t_j^{(0)}$  and  $CW^{(1)}, \dots, CW^{(k)}$ .

4. For  $i$ , the other party's seed  $s_{1-j}^{(i)} \leftarrow \{0, 1\}^\lambda$  is chosen at random,  $t_{1-j}^{(i)} = 1 - t_j^{(i)}$ ,  $\sigma_{1-j}^{(i)} = \sigma_j^{(i)}$ , and  $\tau_{1-j}^{(i)} = \tau_j^{(i)}$ .
5. For  $k \geq i$ : the remaining values  $s_j^{(k+1)} \parallel t_j^{(k+1)}$ ,  $s_{1-j}^{(k+1)} \parallel t_{1-j}^{(k+1)}$ ,  $CW^{(k)}$ ,  $\sigma_j^{(k+1)} \parallel \tau_j^{(k+1)}$ ,  $\sigma_{1-j}^{(k+1)} \parallel \tau_{1-j}^{(k+1)}$ ,  $CW_{leaf}^{(k)}$  all computed honestly, as a function of the previously chosen values.
6. The output of the experiment is  $k_j := \llbracket \alpha \rrbracket_j \parallel s_j^{(1)} \parallel (CW^{(i)})_{i=1..n} \parallel (CW_{leaf}^{(i)})_{i=1..n+1}$ .

$\text{Hyb}_i$  is formally described in Algorithm 10. When  $i = 0$ , the algorithm corresponds to the honest key generation, while when  $i = n + 1$ , it generates a completely random key. We only need to prove that for any  $i \in \{1, \dots, n+1\}$ ,  $\text{Hyb}_{i-1}$  and  $\text{Hyb}_i$  are indistinguishable based on the security of our PRG.

More precisely, let us first consider  $i \leq n$ .

**Claim A.1.** There exists a polynomial  $p'$  such that for any  $(T, \epsilon_{\text{PRG}})$ -secure pseudorandom generator  $G$ , then for every  $i \leq n$ ,  $j \in \{0, 1\}$ , and every non-uniform adversary  $\mathcal{A}$  running in time  $T' \leq T - p'(\lambda)$ , it holds that

$$\left| \Pr[k_j \leftarrow \text{Hyb}_{i-1}(1^\lambda, j); c \leftarrow \mathcal{A}(1^\lambda, k_j) : c = 1] - \right.$$

$$\left. \Pr[k_j \leftarrow \text{Hyb}_i(1^\lambda, j); c \leftarrow \mathcal{A}(1^\lambda, k_j) : c = 1] \right| < \epsilon_{\text{PRG}}$$

*Proof.* Let's fix  $i \in \{1, \dots, n\}$ ,  $j \in \{0, 1\}$ . Let  $\mathcal{A}$  be a  $\text{Hyb}$ -distinguishing adversary with advantage  $\epsilon$  for these values. We use  $\mathcal{A}$  to construct a corresponding PRG adversary  $\mathcal{B}$ . Recall that in the PRG challenge for  $G$ , the adversary  $\mathcal{B}$  is given a value  $r$  that is either computed by sampling a seed  $s \leftarrow \{0, 1\}^\lambda$  and computing  $r = G(s)$ , or is sampled truly at random  $r \leftarrow \{0, 1\}^{2(\lambda+n+2)}$ . Algorithm 11 describes the PRG challenge of  $\mathcal{B}$  embedded in the  $\text{Hyb}$ -distinguishing challenge of  $\mathcal{A}$ .

Now, consider  $\mathcal{B}$ 's success in the PRG challenge as a function of  $\mathcal{A}$ 's success in distinguishing  $\text{Hyb}_{i-1}$  from  $\text{Hyb}_i$ . This means that if  $\mathcal{A}$  succeeds, then  $\mathcal{B}$  will succeed at its challenge, which implies Claim A.1. If, in Algorithm 11,  $r$  is computed *pseudorandomly* using the PRG, then it is clear the generated  $k_j$  is distributed as  $\text{Hyb}_{i-1}(1^\lambda, j)$ .

It remains to show that if  $r$  was sampled at random then the generated  $k_j$  is distributed as  $\text{Hyb}_i(1^\lambda, j)$ . That is, if  $r$  is random, then the corresponding computed values of  $s_{1-j}^{(i+1)}$ ,  $CW^{(i)}$  and  $CW_{leaf}^{(i)}$  are distributed *randomly* conditioned on the values of  $s_j^{(1)} \parallel t_j^{(1)} \parallel (CW^{(i)})_{i=1..i-1} \parallel (CW_{leaf}^{(i)})_{i=1..i-1}$ , and the value of  $t_{1-j}^{(i)}$  is given by  $1 - t_j^{(i)}$ . Note that all re-

maining values (for  $k > i$ ) are computed as a function of the values computed up to step  $i$ .

First, consider  $CW^{(i)}$ , which is computed as such:

$$CW^{(i)} = cw^{(i)} \oplus G(s_j^{(i)}) \oplus r$$

In particular, when  $\alpha[i] = 1$ :

$$\begin{aligned} & cw^{(i)} \oplus r \\ &= ((0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1), (\sigma_0^R \oplus \sigma_1^R \parallel 1, 0^\lambda \parallel 0)) \\ &\oplus ((s_{1-j}^L \parallel t_{1-j}^L, s_{1-j}^R \parallel t_{1-j}^R), \\ &\quad (\sigma_{1-j}^L \parallel \tau_{1-j}^L, \sigma_{1-j}^R \parallel \tau_{1-j}^R)) \\ &= ((0^\lambda \parallel 0, s_j^L \parallel 1), (\sigma_j^R \parallel 1, 0^\lambda \parallel 0)) \\ &\oplus ((s_{1-j}^L \parallel t_{1-j}^L, s_{1-j}^R \oplus s_{1-j}^L \parallel t_{1-j}^R), \\ &\quad (\sigma_{1-j}^L \parallel \tau_{1-j}^L, \sigma_{1-j}^L \oplus \sigma_{1-j}^R \parallel \tau_{1-j}^R)) \end{aligned}$$

When  $\alpha[i] = 0$ :

$$\begin{aligned} & cw^{(i)} \oplus r \\ &= ((s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1)) \\ &\oplus ((s_{1-j}^L \parallel t_{1-j}^L, s_{1-j}^R \parallel t_{1-j}^R), \\ &\quad (\sigma_{1-j}^L \parallel \tau_{1-j}^L, \sigma_{1-j}^R \parallel \tau_{1-j}^R)) \\ &= ((s_j^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, \sigma_j^L \parallel 1)) \\ &\oplus ((s_{1-j}^L \oplus s_{1-j}^R \parallel t_{1-j}^L, s_{1-j}^R \parallel t_{1-j}^R), \\ &\quad (\sigma_{1-j}^L \parallel \tau_{1-j}^L, \sigma_{1-j}^L \oplus \sigma_{1-j}^R \parallel \tau_{1-j}^R)) \end{aligned}$$

Independently of the value of  $\alpha[i]$ , when  $r$  is random, the right hand side of the  $\oplus$  acts as a perfect one-time pad, and so  $CW^{(i)}$  is distributed uniformly.

Consider now  $CW_{leaf}^{(i)}$ , computed as such:

$$CW_{leaf}^{(i)} \leftarrow (-1)^{\tau_1^{(i+1)}} \cdot \left( \sigma_1^{(i+1)} - \sigma_0^{(i+1)} + \alpha[i] \right) \bmod 2^n$$

Since  $\sigma_{1-j}^{(i+1)}$  is distributed randomly conditioned on the previous values computed, it acts as a one-time pad, which ensures that  $CW_{leaf}^{(i)}$  is distributed uniformly.

Now, condition on  $CW^{(i)}$  as well, and consider the value of  $s_{1-j}^{(i+1)} \cdot s_{1-j}^{(i+1)}$  is extracted from  $r \oplus t_{1-j}^{(i)} \cdot CW^{(i)}$ , see Line 15. If  $t_{1-j}^{(i)} = 0$ ,  $s_{1-j}^{(i+1)}$  is immediately uniformly distributed. If  $t_{1-j}^{(i)} = 1$ ,  $r \oplus (t_{1-j}^{(i)} \cdot CW^{(i)}) = r \oplus CW^{(i)} = r \oplus cw^{(i)} \oplus G(s_j^{(i)}) \oplus r = cw^{(i)} \oplus G(s_j^{(i)})$ . When  $\alpha[i] = 1$ , the part of  $cw^{(i)}$  contributing to  $s_{1-j}^{(i+1)}$  is  $s_j^L \oplus s_{1-j}^L$ .  $s_{1-j}^L$  is random and hence acts as a perfect one-time pad, so  $s_{1-j}^{(i+1)}$  is uniformly distributed. When  $\alpha[i] = 0$ , the same result is derived using  $s_{1-j}^R$ .

Finally, consider the value of  $t_{1-j}^{(i+1)}$ . We show that  $t_{1-j}^{(i+1)} = 1 - t_j^{(i)}$ . By construction,

$$t_{1-j}^{(i+1)} = t_{1-j}^{\text{Keep}} \oplus t_{1-j}^{(i)} \cdot t_{CW}^{\text{Keep}}$$

where  $\text{Keep} = \text{L}$  if  $\alpha[i] = 0$  else  $\text{R}$ , and where  $t_{CW}^{\text{Keep}} = 1 \oplus t_0^{\text{Keep}} \oplus t_1^{\text{Keep}}$ . Further more, by noting that  $t_{1-j}^{(i)}$  was

```

Input:  $(1^\lambda, i, j)$ 
Initialisation: Sample random  $\alpha \xleftarrow{\$} \mathbb{Z}_{2^n}$ 
Sample random  $s_0^{(1)}, s_1^{(1)} \xleftarrow{\$} \{0, 1\}^\lambda$  and set  $t_0^{(1)} = 0, t_1^{(1)} = 1$ 
1 for  $k = 1..n$  do
2    $((s_j^L \parallel t_j^L, s_j^R \parallel t_j^R), (\sigma_j^L \parallel \tau_j^L, \sigma_j^R \parallel \tau_j^R)) \leftarrow G(s_j^{(k)})$ 
3   if  $k < i$  then
4      $CW^{(k)} \xleftarrow{\$} \{0, 1\}^{2(\lambda+n+2)}$ 
5   else
6     if  $k = i$  then  $s_{1-j}^{(i)} \xleftarrow{\$} \{0, 1\}^\lambda$  and  $t_{1-j}^{(i)} = 1 - t_j^{(i)}$  ;
7      $((s_{1-j}^L \parallel t_{1-j}^L, s_{1-j}^R \parallel t_{1-j}^R), (\sigma_{1-j}^L \parallel \tau_{1-j}^L, \sigma_{1-j}^R \parallel \tau_{1-j}^R)) \leftarrow G(s_{1-j}^{(k)})$ 
8     if  $\alpha[k]$  then
9        $cw^{(k)} \leftarrow ((0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1), (\sigma_0^R \oplus \sigma_1^R \parallel 1, 0^\lambda \parallel 0))$ 
10      else
11         $cw^{(k)} \leftarrow ((s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, \sigma_0^L \oplus \sigma_1^L \parallel 1))$ 
12       $CW^k \leftarrow cw^{(k)} \oplus G(s_0^{(k)}) \oplus G(s_1^{(k)})$ 
13       $state_{1-j} \leftarrow G(s_{1-j}^{(k)}) \oplus (t_{1-j}^{(k)} \cdot CW^k) = ((state_{1-j,0}, state_{1-j,1}), (state'_{1-j,0}, state'_{1-j,1}))$ 
14      Parse  $s_{1-j}^{(k+1)} \parallel t_{1-j}^{(k+1)} = state_{1-j,\alpha[k]}$  and  $\sigma_{1-j}^{(k+1)} \parallel \tau_{1-j}^{(k+1)} = state'_{1-j,1-\alpha[k]}$ 
15       $state_j \leftarrow G(s_j^{(k)}) \oplus (t_j^{(k)} \cdot CW^k) = ((state_{j,0}, state_{j,1}), (state'_{j,0}, state'_{j,1}))$ 
16      Parse  $s_j^{(k+1)} \parallel t_j^{(k+1)} = state_{j,\alpha[k]}$  and  $\sigma_j^{(k+1)} \parallel \tau_j^{(k+1)} = state'_{j,1-\alpha[k]}$ 
17     if  $k < i$  then
18        $CW_{leaf}^k \xleftarrow{\$} \{0, 1\}^n$ 
19     else
20        $CW_{leaf}^k \leftarrow (-1)^{\tau_1^{(k+1)}} \cdot (\sigma_1^{(k+1)} - \sigma_0^{(k+1)} + \alpha[k]) \bmod 2^n$  if  $k < i$  else  $\{0, 1\}^n$ 
21    $CW_{leaf}^{(n+1)} \leftarrow (-1)^{t_1^{(n+1)}} \cdot (1 - s_0^{(n+1)} + s_1^{(n+1)}) \bmod 2^n$  if  $i \leq n$  else  $\{0, 1\}^n$ 
22 return  $k_j \leftarrow \llbracket \alpha \rrbracket_j \parallel (CW^{(i)})_{i=1..n} \parallel (CW_{leaf}^{(i)})_{i=1..n+1}$ 

```

**Algorithm 10:**  $Hyb_i$ : Hybrid distribution  $i$ , in which the first  $i$  correction words are sampled completely at random, and the remaining correction words are computed honestly.

**Input:**  $(1^\lambda, (i, j), r)$

- 1 Sample random  $\alpha \xleftarrow{\$} \mathbb{Z}_{2^n}$
- 2 Sample  $s_j^{(1)} \xleftarrow{\$} \{0, 1\}^\lambda$  and set  $t_j^{(1)} \leftarrow j$
- 3 **for**  $k = 1..i - 1$  **do**
- 4      $CW^{(k)} \xleftarrow{\$} \{0, 1\}^{2(\lambda+n+2)}$
- 5      $CW_{leaf}^{(k)} \leftarrow \{0, 1\}^n$
- 6      $state_j \leftarrow G(s_j^{(k)}) \oplus (t_j^{(k)} \cdot CW^{(k)}) = ((state_{j,0}, state_{j,1}), (state'_{j,0}, state'_{j,1}))$
- 7     Parse  $s_j^{(k+1)} \parallel t_j^{(k+1)} = state_{j,\alpha[k]}$  and  $\sigma_j^{(k+1)} \parallel \tau_j^{(k+1)} = state'_{j,1-\alpha[k]}$
- 8     Take  $t_{1-j}^{(k+1)} = 1 - t_j^{(k+1)}$
- 9  $((s_j^L \parallel t_j^L, s_j^R \parallel t_j^R), (\sigma_j^L \parallel \tau_j^L, \sigma_j^R \parallel \tau_j^R)) \leftarrow G(s_j^{(i)})$
- 10  $((s_{1-j}^L \parallel t_{1-j}^L, s_{1-j}^R \parallel t_{1-j}^R), (\sigma_{1-j}^L \parallel \tau_{1-j}^L, \sigma_{1-j}^R \parallel \tau_{1-j}^R)) \leftarrow r$  ; // The PRG challenge
- 11 **if**  $\alpha[i]$  **then**  $cw^{(i)} \leftarrow ((0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1), (\sigma_0^R \oplus \sigma_1^R \parallel 1, 0^\lambda \parallel 0))$
- 12 **else**  $cw^{(i)} \leftarrow ((s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, \sigma_0^L \oplus \sigma_1^L \parallel 1))$ ;
- 13  $CW^{(i)} \leftarrow cw^{(i)} \oplus G(s_j^{(i)}) \oplus r$
- 14 **for**  $x = 0, 1$  **do**
- 15      $state_x \leftarrow G(s_x^{(i)}) \oplus (t_x^{(i)} \cdot CW^{(i)})$  **if**  $x = j$  **else**  $r \oplus (t_x^{(i)} \cdot CW^{(i)})$
- 16      $state_x = ((state_{x,0}, state_{x,1}), (state'_{x,0}, state'_{x,1}))$
- 17     Parse  $s_x^{(i+1)} \parallel t_x^{(i+1)} = state_{x,\alpha[i]}$  and  $\sigma_x^{(i+1)} \parallel \tau_x^{(i+1)} = state'_{x,1-\alpha[i]}$
- 18  $CW_{leaf}^{(i)} \leftarrow (-1)^{\tau_1^{(i+1)}} \cdot (\sigma_1^{(i+1)} - \sigma_0^{(i+1)} + \alpha[i]) \bmod 2^n$
- 19 **for**  $k = i + 1..n$  **do**
- 20     **for**  $x = 0, 1$  **do**
- 21          $((s_x^L \parallel t_x^L, s_x^R \parallel t_x^R), (\sigma_x^L \parallel \tau_x^L, \sigma_x^R \parallel \tau_x^R)) \leftarrow G(s_x^{(k)})$
- 22         **if**  $\alpha[k]$  **then**  $cw^{(k)} \leftarrow ((0^\lambda \parallel 0, s_0^L \oplus s_1^L \parallel 1), (\sigma_0^R \oplus \sigma_1^R \parallel 1, 0^\lambda \parallel 0))$
- 23         **else**  $cw^{(k)} \leftarrow ((s_0^R \oplus s_1^R \parallel 1, 0^\lambda \parallel 0), (0^\lambda \parallel 0, \sigma_0^L \oplus \sigma_1^L \parallel 1))$ ;
- 24          $CW^{(k)} \leftarrow cw^{(k)} \oplus G(s_0^{(k)}) \oplus G(s_1^{(k)})$
- 25         **for**  $x = 0, 1$  **do**
- 26              $state_x \leftarrow G(s_x^{(k)}) \oplus (t_x^{(k)} \cdot CW^{(k)}) = ((state_{x,0}, state_{x,1}), (state'_{x,0}, state'_{x,1}))$
- 27             Parse  $s_x^{(k+1)} \parallel t_x^{(k+1)} = state_{x,\alpha[k]}$  and  $\sigma_x^{(k+1)} \parallel \tau_x^{(k+1)} = state'_{x,1-\alpha[k]}$
- 28              $CW_{leaf}^{(k)} \leftarrow (-1)^{\tau_1^{(k+1)}} \cdot (\sigma_1^{(k+1)} - \sigma_0^{(k+1)} + \alpha[k]) \bmod 2^n$
- 29  $CW_{leaf}^{(n+1)} \leftarrow (-1)^{t_1^{(n+1)}} \cdot (1 - s_0^{(n+1)} + s_1^{(n+1)}) \bmod 2^n$
- 30 **return**  $k_j = \llbracket \alpha \rrbracket_j \parallel s_j^{(1)} \parallel (CW^{(i)})_{i=1..n} \parallel (CW_{leaf}^{(i)})_{i=1..n+1}$

**Algorithm 11:** PRG Challenge for adversary  $\mathcal{B}$

set to  $1 - t_j^{(i)}$ ,

$$\begin{aligned}
 & t_j^{(i+1)} \oplus t_{1-j}^{(i+1)} \\
 &= (t_j^{\text{Keep}} \oplus t_j^{(i)} \cdot t_{CW}^{\text{Keep}}) \oplus (t_{1-j}^{\text{Keep}} \oplus t_{1-j}^{(i)} \cdot t_{CW}^{\text{Keep}}) \\
 &= t_j^{\text{Keep}} \oplus t_{1-j}^{\text{Keep}} \oplus (t_j^{(i)} \oplus t_{1-j}^{(i)}) \cdot t_{CW}^{\text{Keep}} \\
 &= t_j^{\text{Keep}} \oplus t_{1-j}^{\text{Keep}} \oplus 1 \cdot (1 \oplus t_0^{\text{Keep}} \oplus t_1^{\text{Keep}}) \\
 &= 1
 \end{aligned}$$

Combining these pieces, we have that in the case of a random PRG challenge  $r$ , the resulting distribution of  $k_j$  as generated by  $\mathcal{B}$  is precisely distributed as is  $\text{Hyb}_i(1^\lambda, j)$ . Thus, the advantage of  $\mathcal{B}$  in the PRG challenge experiment is equivalent to the advantage  $\epsilon$  of  $\mathcal{A}$  in distinguishing  $\text{Hyb}_{i-1}(1^\lambda, j)$  from  $\text{Hyb}_i(1^\lambda, j)$ . The runtime of  $\mathcal{B}$  is equal to the runtime of  $\mathcal{A}$  plus a fixed polynomial  $p'(\lambda)$ . Thus for any  $T' \leq T - p'(\lambda)$ , it must be that the distinguishing advantage  $\epsilon$  of  $\mathcal{A}$  is bounded by  $\epsilon_{\text{PRG}}$ , which concludes the proof of Claim A.1.

Combining all the steps, for  $i \in \{1, \dots, n\}$ , we have thus proven that  $\text{Hyb}_0(1^\lambda, j)$  and  $\text{Hyb}_n(1^\lambda, j)$  are computationally indistinguishable for any adversary, for  $j \in \{0, 1\}$ . On the other hand, we can also prove:

**Claim A.2.**

$$\text{Hyb}_n(1^\lambda, j) = \text{Hyb}_{n+1}(1^\lambda, j)$$

*Proof.* In  $\text{Hyb}_{n+1}$ ,  $CW_{\text{leaf}}^{(n+1)} \stackrel{\$}{\leftarrow} \{0, 1\}^n$ . In  $\text{Hyb}_n$ ,  $CW_{\text{leaf}}^{(n+1)} = (-1)^{t_1^{(n+1)}} \cdot (1 - s_0^{(n+1)} + s_1^{(n+1)}) \bmod 2^n$ , with  $s_{1-j}^{(n+1)}$  distributed randomly conditioned on the previous values computed.  $s_{1-j}^{(n+1)}$  acts as a one-time pad which perfectly hides other values. Hence,  $CW_{\text{leaf}}^{(n+1)}$  is also uniformly distributed in this case.

This concludes the proof of security of our FSS comparison protocol.

## B Encoding Precision

We have studied the impact of lowering the encoding space of the input to our function secret sharing protocol from  $\mathbb{Z}_{2^{32}}$  to  $\mathbb{Z}_{2^k}$  with  $k < 32$ . Finding the lowest  $k$  guaranteeing good performance is an interesting challenge as the function keys size is directly proportional to it. This has to be done together with reducing fixed precision from 3 decimals down to 1 decimal to ensure private values aren't too big, which would result in higher failure rate in our private comparison protocol. We have reported in Table 8 our findings on Network-1, which is pre-trained and then evaluated in a private fashion.

Decimals	$\mathbb{Z}_{2^{12}}$	$\mathbb{Z}_{2^{16}}$	$\mathbb{Z}_{2^{20}}$	$\mathbb{Z}_{2^{24}}$	$\mathbb{Z}_{2^{28}}$	$\mathbb{Z}_{2^{32}}$
1	-	-	-	-	-	9.5
2	69.4	96.0	97.9	98.1	98.0	98.1
3	10.4	76.2	96.9	98.1	98.2	98.1
4	9.7	14.3	83.5	97.4	98.1	98.2

**Table 8.** Accuracy (in %) of Network-1 given different precision and encoding spaces

What we observe is that 3 decimals of precision is the most appropriate setting to have an optimal precision while allowing to slightly reduce the encoding space down to  $\mathbb{Z}_{2^{24}}$  or  $\mathbb{Z}_{2^{28}}$ . Because this is not a massive gain and in order to keep the failure rate in comparison very low, we have kept  $\mathbb{Z}_{2^{32}}$  for all our experiments.


## C Implementation Details

### C.1 Pseudo-Random Generator

The PRG is implemented using a Matyas-Meyer-Oseas one-way compression function as in [61], with an AES block cipher. We concatenate several fixed key block ciphers to achieve the desired output length:  $G(x) = E_{k_1}(x) \oplus x \parallel E_{k_2}(x) \oplus x \parallel \dots$ . Those keys are fixed and hard-coded. We set  $\lambda = 127$  to be able to use only 2 blocks for equality and 4 blocks for comparison. Note that for comparison we would theoretically only need 3 blocks, although our current implementation uses 4.

### C.2 Unrolling Convolutions

Figure 3 illustrates how to transform a convolution operation into a single matrix multiplication.



**Fig. 3.** Illustration of unrolling a convolution with kernel size  $k = 2$  and stride  $s = 2$ .

### C.3 MaxPool and Optimisation

Figure 4 illustrates how MaxPool uses ideas from matrix unrolling and argmax computation. Notations present in the figure are consistent with the explanation of argmax using pairwise comparison in Section 4.2. The  $m \times m$  matrix is first unrolled to a roughly  $(m/s)^2 \times k^2$  matrix. It is then expanded on  $k^2$  layers, each of which is shifted by a step of 1. Next,  $(m/s)^2 k^2 (k^2 - 1)$  pairwise comparisons are then applied simultaneously between the first layer and the other ones, and for each  $x_i$  we sum the result of its  $k - 1$  comparison and check if it equals  $k - 1$ . We multiply this boolean by  $x_i$  and sum up along a line (like  $x_0$  to  $x_3$  in the figure). Last, we restructure the matrix back to its initial structure.




Fig. 4. Illustration of MaxPool with kernel size  $k = 2$  and stride  $s = 2$ .

In addition, when the kernel size  $k$  is 2, rows are only of length 4 and it can be more efficient to use a binary tree approach instead, i.e. compute the maximum of columns 0 and 1, 2 and 3 and the max of the result: it requires  $2 \log_2(k^2) = 4$  rounds of communication but only approximately  $(k^2 - 1)(m/s)^2$  private comparisons, compared to a fixed 3 rounds and approximately  $k^4(m/s)^2$ . We found in practice that this  $4 \times$  speed-up factor in computation is worth an additional communication round.

Interestingly, average pooling can be computed locally on the shares without interaction because it only includes mean operations, but we didn't replace MaxPool operations with average pooling to avoid distorting existing neural networks architecture.

### C.4 Breaking Ties in Argmax

Algorithm 12 provides a probabilistic method to break ties in the argmax output, which can be used on secret shared input as well.

**Input:**  $\vec{\delta} = (\delta_1, \dots, \delta_m)$ , with  $\forall i, \delta_i \in \{0, 1\}$   
**Output:**  $\delta_j$  with  $j \stackrel{\$}{\leftarrow} \{i, \delta_i = 1\}$

- 1  $\vec{x} = \text{CumSum}(\vec{\delta})$
- 2  $r \stackrel{\$}{\leftarrow} [0, \vec{x}[-1][$
- 3  $\vec{c} = \vec{x} > r$
- 4 Compute  $\vec{c}_{\gg}$  by shifting  $\vec{c}$  by one number on the right and padding with 0
- 5 **return**  $\vec{c} - \vec{c}_{\gg}$

Algorithm 12: BreakTies algorithm to guarantee one-hot output

### C.5 BatchNorm Approximation

The BatchNorm layer is the only one in our implementation which requires a polynomial approximation during training. We have therefore experimented how this approximation can alter the behaviour of a deep network such as ResNet18 on a simple dataset like the hymenoptera dataset<sup>3</sup>. We follow the PyTorch transfer learning tutorial<sup>4</sup> and use a pretrained version of ResNet18. We retrain all the layers for 25 epochs, but we replace the BatchNorm layers with our approximated version which behaves as such: as a general rule the inverse for the variance is computed using 3 iterations of the Newton methods and we use the result of the previous batch as an initial approximation. For the first batch of the epoch where no approximation is available, we use instead 50 iterations. For the first BatchNorm layer, since the variance can change significantly because of the input diversity, we systematically use more iterations, up to 60. For the last layer, as the variance is very small, its inverse can have a large amplitude. Therefore we don't use the result of the previous batch and perform 10 iterations instead. We report in Figure 5 the evolution of the error as the model train on more batches. As one can see, the error dramatically shrinks after 10 batches, which shows how beneficial it is to use previous batch computations. It almost stays below 5%, except for the error of the first layer, for which we would need even more iterations to have significant improvements. The average error per layer is reported in Figure 6, and is around 2%.

This is of course an experimental result, but it shows that the BatchNorm approximation can be tailored for

<sup>3</sup> [https://download.pytorch.org/tutorial/hymenoptera\\_data.zip](https://download.pytorch.org/tutorial/hymenoptera_data.zip)

<sup>4</sup> [https://pytorch.org/tutorials/beginner/transfer\\_learning\\_tutorial.html](https://pytorch.org/tutorials/beginner/transfer_learning_tutorial.html)






Fig. 5. BatchNorm relative error as training goes on




Fig. 6. Relative error of the BatchNorm layers across the model

specific models to allow for efficient training with little round overhead. Table 1 shows that we can indeed achieve a high accuracy, especially using the initial guess of the previous batch, although it remains below the accuracy of a model trained with an exact BatchNorm.

## D Extended Results about Private Inference

We provide additional results about our inference experiment in Table 9.

## E Datasets and Networks Architecture

### E.1 Datasets

*This section is taken almost verbatim from [60].*

We select 4 datasets popularly used for training image classification models: MNIST [43], CIFAR-10 [39],

64×64 Tiny Imagenet [62] and Hymenoptera, a subset of the Imagenet dataset [53] composed 224×224 pixel images.

**MNIST** MNIST [43] is a collection of handwritten digits dataset. It consists of 60,000 images in the training set and 10,000 in the test set. Each image is a 28×28 pixel image of a handwritten digit along with a label between 0 and 9. We evaluate Network-1, Network-2, and the LeNet network on this dataset.

**CIFAR-10** CIFAR-10 [39] consists of 50,000 images in the training set and 10,000 in the test set. It is composed of 10 different classes (such as airplanes, dogs, horses, etc.) and there are 6,000 images of each class with each image consisting of a colored 32×32 image. We perform private training of AlexNet and inference of VGG16 on this dataset.

**Tiny ImageNet** Tiny ImageNet [62] consists of two datasets of 100,000 training samples and 10,000 test samples with 200 different classes. The first dataset is composed of colored 64×64 images and we use it with AlexNet and VGG16. The second is composed of colored 224×224 images and is used with ResNet18.

**Hymenoptera** Hymenoptera is a dataset extracted from the ImageNet database. It is composed of 245 training and 153 test colored 224×224 images, and was first proposed as a transfer learning task.

### E.2 Model Description

We have selected 6 models for our experimentations. Description on the first 5 models is taken verbatim from [60].

**Network-1** A 3-layered fully-connected network with ReLU used in SecureML [46].

**Network-2** A 4-layered network selected in MiniONN [44] with 2 convolutional and 2 fully-connected layers, which uses MaxPool in addition to ReLU activation.

**LeNet** This network, first proposed by LeCun et al. [42], was used in automated detection of zip codes and digit recognition. The network contains 2 convolutional layers and 2 fully connected layers.

**AlexNet** AlexNet is the famous winner of the 2012 ImageNet ILSVRC-2012 competition [40]. It has 5 convolutional layers and 3 fully connected layers and it can use batch normalization layers for stability and efficient training.

**Table 9.** Comparison of the inference time between secure frameworks on several popular neural network architectures. For each network we report in order the batch size used and the size of the preprocessing material in MB for the CPU setting.

Framework Dataset		Network-1			Network-2			LeNet		
		Batch Size		Preprocessing	Batch Size		Preprocessing	Batch Size		Preprocessing
		CPU	GPU	Comm. (MB)	CPU	GPU	Comm. (MB)	CPU	GPU	Comm. (MB)
AriaNN	MNIST	128	128	0.36	128	128	9.98	128	128	14.7

Framework Dataset		AlexNet			VGG16			ResNet18		
		Batch Size		Preprocessing	Batch Size		Preprocessing	Batch Size		Preprocessing
		CPU	GPU	Comm. (MB)	CPU	GPU	Comm. (MB)	CPU	GPU	Comm. (MB)
AriaNN	CIFAR-10	128	128	24.6	64	14	277	-	-	-
AriaNN	64×64 ImageNet	128	64	88.4	16	3	1124	-	-	-
AriaNN	224×224 Hymenoptera	-	-	-	-	-	-	8	1	3254

**VGG16** VGG16 is the runner-up of the ILSVRC-2014 competition [58]. VGG16 has 16 layers and has about 138M parameters.

**ResNet18** ResNet18 [30] is the runner-up of the ILSVRC-2015 competition. It is a convolutional neural network that is 18 layers deep, and has 11.7M parameters. It uses batch normalisation and we are the first private deep learning framework to evaluate this network.

### E.3 Models Architecture

Unless otherwise specified, our models follow their standard architecture as provided by the torchvision library (version 0.5), except for smaller models such as Network-1, Network-2 and LeNet which are respectively detailed in [46], [44] and [42].

For the CIFAR-10 version of AlexNet, we follow the architecture of [60] which includes BatchNorm layers and is available on the FALCON GitHub. For the 64×64 Tiny Imagenet version of AlexNet however, we used the standard architecture from PyTorch since it allows us to have a pretrained network and the version of FALCON seemed non-standard.

We have adapted the classifier parts of AlexNet, VGG16 and ResNet18 to the different datasets we use.

- For AlexNet on Tiny Imagenet, we use 3 fully connected layers with respectively 1024, 1024 and 200 neurons, and ReLU activations between them.
- For VGG16 we use 3 fully connected layers with respectively 4096, 4096 and 10 or 200 neurons for the last layer depending on the dataset used. We also use ReLU activations between them.

- For ResNet18, we use a single fully connected layer to map the 512 output logits to the appropriate number of classes.

Note also that we permute ReLU and Maxpool where applicable like in [60], as this is strictly equivalent in terms of output for the network and reduces the number of comparisons. More generally, we don't proceed to any alteration of the network behaviour except with the approximation on BatchNorm. This improves the usability of our framework as it allows us to use pre-trained neural networks from standard deep learning libraries like PyTorch and to encrypt them with a single line of code.