Love or Hate? Share or Split? Privacy-Preserving Training Using Split Learning and Homomorphic Encryption

1st Tanveer Khan
Tampere University
Tampere, Finland
tanveer.khan@tuni.fi

2nd Khoa Nguyen
Tampere University
Tampere, Finland
khoa.nguyen@tuni.fi

3rd Antonis Michalas
Tampere University, Finland,
University of Westminster
antonios.michalas@tuni.fi

4th Alexandros Bakas
Tampere University and Nokia Bell Labs
Tampere, Finland
alexandros.bakas@tuni.fi

***Abstract*—Split learning (SL) is a new collaborative learning technique that allows participants, e.g. a client and a server, to train machine learning models without the client sharing raw data. In this setting, the client initially applies its part of the machine learning model on the raw data to generate activation maps and then sends them to the server to continue the training process. Previous works in the field demonstrated that reconstructing activation maps could result in privacy leakage of client data. In addition to that, existing mitigation techniques that overcome the privacy leakage of SL prove to be significantly worse in terms of accuracy. In this paper, we improve upon previous works by constructing a protocol based on U-shaped SL that can operate on homomorphically encrypted data. More precisely, in our approach, the client applies homomorphic encryption on the activation maps before sending them to the server, thus protecting user privacy. This is an important improvement that reduces privacy leakage in comparison to other SL-based works. Finally, our results show that, with the optimum set of parameters, training with HE data in the U-shaped SL setting only reduces accuracy by 2.65% compared to training on plaintext. In addition, raw training data privacy is preserved.**

***Index Terms*—Activation Maps, Homomorphic Encryption, Machine Learning, Privacy, Split Learning**

I. Introduction

SL [1] is a collaborative learning approach that divides a Machine Learning (ML) into two parts: the client-side and the server-side. It facilitates training the Deep Neural Networks (DNN) among multiple data sources, while mitigating the need to directly share raw labeled data with collaboration parties. The advantages of SL are multifold: *(i)* it allows multiple parties to collaboratively train a neural network, *(ii)* it allows users to train ML models without sharing their raw data with a server running part of a DNN model, thus preserving user privacy, *(iii)* it protects both the client and the server from revealing their parts of the model, and *(iv)* it reduces the client's computational overhead by not running the entire model (i.e. utilizing a smaller number of layers) [2].

Though SL offers an extra layer of privacy protection by definition, there are no works exploring how it is combined with popular techniques that promise to preserve user privacy (e.g. encryption). In [3], the authors studied whether SL can handle sensitive time-series data and demonstrated that SL alone is *insufficient* when performing privacy-preserving training for 1-dimensional (1D) CNN models. More precisely, the authors showed raw data can be reconstructed from the activation maps of the intermediate split layer. The authors also employed two mitigation techniques, adding hidden layers and applying differential privacy to reduce privacy leakage. However, based on the results, none of these techniques can effectively reduce privacy leakage from all channels of the SL activation. Furthermore, both these techniques result in significantly reducing the joint model's accuracy.

In this paper, we focus on training an ML model in a privacy-preserving manner, where a client and a server collaborate to train the model. More specifically, we construct a model that uses Homomorphic Encryption (HE) [4] to mitigate privacy leakage in SL. In our model, the client first encrypts the activation maps and then sends the encrypted activation maps to the server. The encrypted activation maps do *not* reveal anything about the raw data (i.e. it is *not* possible to reconstruct the original raw data from the encrypted activation maps).

Contributions: The main contributions of this paper are:

- **C1.** We construct a U-shaped split 1D CNN model and experiment using plaintext activation maps sent from the client to the server. Through the U-shaped 1D CNN model, clients do *not* need to share the ground truth labels with the server – this is an important improvement that reduces privacy leakage compared to [3].
- **C2.** We constructed the HE version of the U-shaped SL technique. In the encrypted U-shaped SL model, the client encrypts the activation map using HE and sends it to the server. The advantage of HE encrypted U-shaped SL over the plaintext U-shaped SL is that server performs computation over encrypted activation maps.
- **C3.** To assess the applicability of our framework, we performed experiments on the PTB-XL heartbeat datasets [5], currently being the largest open-source electrocardiography dataset to our knowledge. We experimented with activation maps of lengths 256 for both plaintext and homomorphically encrypted activation maps and we have measured the model's performance by considering training duration, test accuracy, and communication cost.

Organization: The rest of the paper is organized as follows: In section II, we provide necessary background for 1D CNN, HE and SL. In section III, we present works in the area of SL, followed by design and implementation of the split 1D CNN

training protocols in section IV, extensive experimental results in section V and conclude the paper in section VI.

II. Preliminaries

A. Convolutional Neural Network

CNN is a special type of neural network using convolution layers to extract features from data [6]. In this work, we employ a 1D CNN [3], [7] as a feature detector and classi er for a heartbeat datasets, namely PTB-XL [5]. The employed 1D CNN consists of the following layers stacked on top of each other, namely: Conv1D, Leaky ReLU, max pooling and a single fully connected and a softmax layer (see Figure 1).

B. Homomorphic Encryption

HE is an emerging cryptographic technique for computations on encrypted data. HE schemes are divided into three main categories according to their functionality: Partial HE [8], leveled (or somewhat) HE [4], and fully HE [9]. Each scheme has its own bene ts and disadvantages. In this work, we use the CKKS Leveled HE scheme [4]. CKKS allows users to do additions and a limited number of multiplications on vectors of complex values (and hence, real values too). The most important parameters of the CKKS scheme are Polynomial Modulus P, Coef cient Modulus C and Scaling Factor Δ.

C. Split Learning

Although the con guration of SL could potentially take many forms, widely known con gurations are simple vanilla split and U-shaped SL. In a simple two-party vanilla split, a client and a server collaborate to train the split model with no access to each other's parts. More speci cally, the layers of a DNN model are split into two parts (A and B) and distributed to the client and the server. The client owning the data, uses forward propagation to train part A (comprising the rst few layers) and sends the activated output from the split layer (part A's nal layer) to the server. The server performs the forward training on the outputs activated from the client using part B. Furthermore, the server performs the backward propagation on part B, only returning to the client the gradients of the split layer's active outputs (part B's rst layer) to complete the backward propagation on part A. This process is repeated until the model converges. Although the client and server do not share any raw input data, this con guration requires label sharing. The sharing of labels can be eliminated by using a U-shaped con guration. The U-shaped SL con guration is different from the simple vanilla setup, as it does not require of clients to share labels [2]. On the server-side, the network is wrapped at the end layer, and the outputs are sent back to client entities (see Figure 1). Clients generate gradients from the end layers and utilize them for backward propagation without revealing the corresponding labels.

III. Related Work

Initially, it was believed that SL is a promising approach in terms of client raw data protection, and for doing PPML [10], as it does not share raw data and also has the bene t of

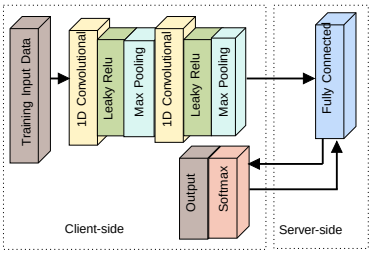

Fig. 1: U-shaped Split-Learning

disclosing the model's architecture and weights.For example, Gupta and Raskar [1] predicted that reconstructing raw data on the client-side, while using SL would be dif cult. In addition, the authors of [2] employed the SL model to the healthcare applications to protect the users' personal data. Vepakomma et al. found that SL outperforms other collaborative learning techniques like federated learning [11] in terms of accuracy [2]. However, SL provides data privacy on the grounds that only activation maps are shared between the parties. Different studies showed the possibility of privacy leakage in SL. In [12], the authors analyzed the privacy leakage of SL and found a considerable leakage from the split layer in the 2D CNN model. Furthermore, the authors mentioned that it is possible to reduce the distance correlation (a measure of dependence) between the split layer and raw data by slightly scaling the weights of all layers before the split. This works well in models with a large number of hidden layers before the split.

The work of Abuadbba et al. [3] is the rst study exploring whether SL can deal with time-series data. According to the results, SL can be applied to a model without the model accuracy degradation. The authors also showed that when SL is directly adopted into 1D CNN models for time series data could result in signi cant privacy leakage. Two mitigation techniques were employed to limit the potential privacy leakage in SL: (i) increasing the number of layers before the split on the client-side and (ii) applying differential privacy to the split layer activation before sending the activation map to the server. However, both techniques suffer from a loss of model accuracy, particularly when differential privacy is used.

In [3], during the forward propagation, the client sends the activation map in plaintext to the server, while the server can easily reconstruct the original raw data from the activated vector of the split layer leading to clear privacy leakage. In our work, we constructed a training protocol, where, instead of sending plaintext activation maps, the client rst conducts an encryption using HE and then sends said maps to the server. In this way, the server is unable to reconstruct the original raw data, but can still perform a computation on the encrypted activation maps, enabling the the training process.

IV. Split Model Training Protocols

In this section, we rst present the protocol for training the U-shaped split 1D CNN on plaintext activation maps, followed by the protocol for training the U-shaped split 1D CNN on encrypted activation maps.

A. Actors in the U-Split Learning Model

We have two parties: client and a server. Each party plays a specific role and has access to certain parameters. We construct the U-shaped split 1D CNN in such a way that the first few as well as the last layer are on the client-side, while the remaining layers are on the server-side, as demonstrated in Figure 1. The client and server collaborate to train the split model by sharing the activation maps and gradients. To elaborate further, we describe their respective roles and access as follows:

Client: In the plaintext version, the client holds two Conv1D layers and can access their weights and biases in plaintext. However, other layers such as Max Pooling, Leaky ReLU, Softmax do not have associated weights and biases. Apart from these, in the HE encrypted version, the client is also responsible for generating the context for HE and has access to all context parameters including Polynomial modulus P, Coefficient modulus C, Scaling factor $$, Public key pk and Secret key sk. Note that for both training on plaintext and encrypted activation maps, the raw data examples x's and their corresponding labels y's reside on the client side and are never sent to the server during the training process.

Server: The computation performed on the server-side is limited to only one linear layer. The reason for having only one linear layer on the server-side is due to computational constraints when training on HE encrypted data. Hence, the server can exclusively access the weights and biases of this linear layer. Regarding the HE context parameters, the server has access to P, C, $$, and pk shared by the client, with the exception of the sk. Not holding the secret key, the server cannot decrypt the HE encrypted activation maps sent from the client. The hyperparameters shared between the client and the server include the learning rate ($$), batch size (n), number of batches to be trained (N), and the number of training epochs (E).

B. Plaintext Activation Maps

We have used algorithm 1 and algorithm 2 to train the U-shaped split 1D CNN model. First, the client and server start the socket initialization process and synchronize the hyperparameters $; n; N; E$. They also initialize the weights of their layers using the randomly initialized weights

During the forward propagation phase, the client forward-propagates the input x until the l^{th} layer and sends the activation $a^{(l)}$ to the server. The server continues to forward-propagate and sends the output $a^{(L)}$ to the client. Next, the client applies the Softmax function on $a^{(L)}$ to get $\hat{y}$ and calculates the error J using the loss function $J = L(\hat{y}; y)$.

The client starts the backward propagation by calculating and sending the gradient of the error w.r.t $a^{(L)}$, i.e. $\frac{\partial J}{\partial a^{(L)}}$, to the server. The server continues the backward propagation, calculates $\frac{\partial J}{\partial a^{(l)}}$ and sends $\frac{\partial J}{\partial a^{(l)}}$ to the client. After receiving the gradients $\frac{\partial J}{\partial a^{(l)}}$ from the server, the backward propagation continues to the first hidden layer on the client-side. Note that the exchange of information between client and server in these algorithms takes place in plaintext. As can be seen in algorithm 1, the client sends the activation maps $a^{(l)}$ to the server in plaintext and receives the output of the linear layer $a^{(L)}$ from the server in plaintext. The same applies on the server side: receiving $a^{(l)}$ and sending $a^{(L)}$ in the plaintext as can be seen in algorithm 2. Shaul et al. [3] showed that the exchange of plaintext activation maps between client and server using SL reveals important information regarding the client's raw sequential data. Later, in subsection IV-E we show in detail how passing the forward activation maps from the client to the server in the plaintext will result in information leakage. To mitigate this privacy leakage, we propose the protocol, where the client encrypts the activation maps before sending them to the server, as described in subsection IV-C.

Algorithm 1: Client Side

```
Initialization:
    s    socket initialized with port and address;
    s.connect
    ; n; N; E    s:synchronize ()
    f w^(i); b^(i) g_{8 i 2f 0::l g}    initialize using
    f z^(i) g_{8 i 2f 0::l g}; f a^(i) g_{8 i 2f 0::l g} ;
    [ @J/@z^(i)  8 i 2f 0::l g ;  @J/@a^(i)  8 i 2f 0::l g ;
    for e 2 E do
        for each batch (x; y) generated from D do
            Forward propagation :
                O:zero _grad ()
                a^0    x
                for i  1 to l do
                    z^(i)    f^(i) a^(i 1)
                    a^(i)    g^(i) z^(i)
                end
                s:send (a^(l))
                s:receive (a^(L))
                y^    Softmax a^(L)
                J  L (y^; y)
            Backward propagation :
                Compute  @J/@y^  &  @J/@a^(L)
                s:send  @J/@a^(L)
                s:receive  @J/@a^(l)
                for i  1 to l do                    o
                    Compute  @J/@w^(i) ; @J/@b^(i)
                        Update w^(i); b^(i)
                end
        end
end
```

C. Encrypted Activation Maps

The protocol for training the U-shaped 1D CNN with a homomorphically encrypted activation map consists of four phases: initialization, forward propagation, classification, and backward propagation. The initialization phase only takes place once at the beginning of the procedure, whereas the other phases continue until the model iterates through all epochs. Each phase is described in detail in the following subsections.

Initialization: The initialization phase consists of socket initialization, context generation, and random weight loading. The client first establishes a socket connection to the server and synchronizes the four hyperparameters $; n; N; E$ with the server, shown in algorithm 3 and algorithm 4. These parameters

Algorithm 2: Server Side

```
Initialization:
s    socket initialized with port and address;
s.connect
 ; n; N; E      s:synchronize   ()
f w^(i); b^(i) g_{8 i 2f 0::l g}      initialize using
f z^(i) g_{8 i 2f l+1 ::L g}   ;
  @J/@z^(i)                    ;
                 8 i 2f l+1 ::L g
for e 2 E do
    for i  1 to N do
        Forward propagation      :
            O:zero _grad ()
            s:receive (a^(l))
            a^(L)      f^(i) a^(l)
            s:send  a^(L)
        Backward propagation      :
            s:receive  @J/@a^(L)
            Compute  @J/@w^(L) ; @J/@b^(L)
            Updatew^(L); b^(L)
            Compute @J/@a^(l)
            s:send  @J/@a^(l)
    end
end
```

Algorithm 3: Client Side

```
Context Initialization:
    ctx_pri;   P  ; C;   ; pk; sk
    ctx_pub;   P  ; C;   ; pk
    s:send (ctx_pub)
for e in E do
    for  each batch(x; y) generated fromD  do
        Forward propagation      :
            O:zero _grad ()
            a^0      x
            for i      1 to l do
                z^(i)      f^(i)  a^(i 1)
                a^i      g^(i)  z^(i)
            end
            a^(l)      HE:Enc  pk; a^(l)
            s:send  (a^(l))
            s:receive  (a^(L))
            a^(L)      HE:Dec  sk; a^(L)
            ŷ      Softmax  a^(L)
            J  L  (ŷ; y)
            Backward propagation      :
            Compute  @J/@ŷ & @J/@a^(L) & @J/@w^(L)
            s:send  @J/@a^(L) & @J/@w^(L)
            s:receive  @J/@a^(l)
            for i    l down to1 do
                Compute  @J/@w^(i) ; @J/@b^(i)
                Updatew^(i); b^(i)
            end
    end
end
```

must be synchronized on both sides to be trained in the same way. Also, the weights on the client and server are initialized with the same set of corresponding weights in the local model to accurately assess and compare the influence of SL on performance. On both the client and the server sides, $w^{(i)}$, are initialized using corresponding parts of. The activation map at layer i ($a^{(i)}$), output tensor of Conv1D layer $z^{(i)}$), and the gradients are initially set to zero. In this phase, the context generated is a specific object that holds pk and sk of the HE scheme as well as certain other parameters like P, C and .

Further information on the HE parameters and how to choose the best-suited parameters can be found in the TenSEAL's benchmarks tutorial. As shown in algorithm 3 and algorithm 4, the context is either public (ctx_{pub}) or private (ctx_{pri}) depending on whether it holds the secret key sk. Both the ctx_{pub} and ctx_{pri} have the same parameters, though ctx_{pri} holds sk and ctx_{pub} does not. The server does not have access to sk as the client only shares the ctx_{pub} with the server. After completing the initialization phase, both the client and server proceed to the forward and backward propagation phases.

Forward propagation: The forward propagation starts on the client side. The client first zeroes out the gradients for the batch of data $(x; y)$. He then begins calculating $a^{(l)}$ from x, algorithm 3 where each $f^{(i)}$ is a Conv1D layer.

The Conv1D layer can be described as following: given a 1D input signal that contains C channels, where each channel $x_{(i)}$ is a 1D array $i(2 f 1; : : : ; C g)$, a Conv1D layer produces an output that contains C^0 channels. The j^{th} output channel $y_{(j)}$, where $j \; 2 f 1; : : : ; C^0 g$, can be described as

$$y_{(j)} = b_{(j)} + \sum_{i=1}^{C} w_{(i)} ? x_{(i)}; \qquad (1)$$

where $w_{(i)}; i \; 2 f 1; : : : ; C g$ are the weights, $b_{(j)}$ are the biases of the Conv1D layer, and ? is the 1D cross-correlation operation. The ? operation can be described as

$$z(i) = (w ? x)(i) = \sum_{j=0}^{m \; 1} w(j) \; x(i + j); \qquad (2)$$

where $z(i)$ denotes the i^{th} element of the output vector z, and i starts at 0. Here, the size of the 1D weighted kernel is .

In algorithm 3, $g^{(i)}$ can be seen as the combination of Max Pooling and Leaky ReLU functions. The final output activation maps of the l^{th} layer from the client is $a^{(l)}$. The client then homomorphically encrypts $a^{(l)}$ and sends the encrypted activation maps $\overline{a^{(l)}}$ to the server. In algorithm 4, the server receives $\overline{a^{(l)}}$ and then performs forward propagation, which is a linear layer evaluated on HE encrypted data $\overline{a^{(l)}}$ as

$$\overline{a^{(L)}} = \overline{a^{(l)}} w^{(L)} + b^{(L)}: \qquad (3)$$

After that, the server sends $\overline{a^{(L)}}$ to the client (algorithm 4). Upon reception, the client decrypts $\overline{a^{(L)}}$ to get $a^{(L)}$, performs Softmax on $a^{(L)}$ to produce the predicted output $\hat{y}$ and calculate the loss J, as can be seen in algorithm 3. Having finished the forward propagation we may move on to the backward propagation part of the protocol.

Backward propagation: After calculating J, the client starts the backward propagation by initially computing $\frac{@J}{@\hat{y}}$ and then $\frac{@J}{@a^{(L)}}$ and $\frac{@J}{@w^{(L)}}$ using the chain rule (algorithm 3).

$$\frac{\partial J}{\partial a^{(L)}} = \frac{\partial J}{\partial \hat{y}} \frac{\partial \hat{y}}{\partial a^{(L)}}; \text{and} \quad \frac{\partial J}{\partial w^{(L)}} = \frac{\partial J}{\partial a^{(L)}} \frac{\partial a^{(L)}}{\partial w^{(L)}} \tag{4}$$

Following, the client sends $\frac{\partial J}{\partial a^{(L)}}$ and $\frac{\partial J}{\partial w^{(L)}}$ to the server. Upon reception, the server computes $\frac{\partial J}{\partial b^{(L)}}$ by simply doing $\frac{\partial J}{\partial b^{(L)}} = \frac{\partial J}{\partial a^{(L)}}$, based on equation (3). The server then updates the weights and biases of his linear layer according to equation (5).

$$w^{(L)} = w^{(L)} - \frac{\partial J}{\partial w^{(L)}}; \quad b^{(L)} = b^{(L)} - \frac{\partial J}{\partial b^{(L)}} \tag{5}$$

Next, the server calculates

$$\frac{\partial J}{\partial a^{(l)}} = \frac{\partial J}{\partial a^{(L)}} \frac{\partial a^{(L)}}{\partial a^{(l)}}; \tag{6}$$

and sends $\frac{\partial J}{\partial a^{(l)}}$ to the client. After receiving $\frac{\partial J}{\partial a^{(l)}}$, the client calculates gradients of J w.r.t the weights and biases of the Conv1D layer using the chain-rule: as:

$$\frac{\partial J}{\partial w^{(i-1)}} = \frac{\partial J}{\partial w^{(i)}} \frac{\partial w^{(i)}}{\partial w^{(i-1)}} \tag{7}$$

$$\frac{\partial J}{\partial b^{(i-1)}} = \frac{\partial J}{\partial b^{(i)}} \frac{\partial b^{(i)}}{\partial b^{(i-1)}} \tag{8}$$

After calculating the gradients $\frac{\partial J}{\partial w^{(i)}}$; $\frac{\partial J}{\partial b^{(i)}}$, the client updates $w^{(i)}$ and $b^{(i)}$ using the Adam optimization algorithm [13].

Algorithm 4: Server Side

```
Context Initialization:
    s:receive (ctx_pub)
for e in E do
    for i 1 to N do
        Forward propagation :
            O:zero_grad()
            s:receive (a^(l))
            a^(L)  HE:Eval f^(i) a^(l)
            s:send a^(L)
        Backward propagation :
            s:receive ∂J/∂a^(L) & ∂J/∂w^(L)
            Compute ∂J/∂b^(L)
            Update w^(L); b^(L)
            Compute ∂J/∂a^(l)
            s:send ∂J/∂a^(l)
    end
end
```

V. Performance Analysis

A. ECG Datasets

We evaluate our method on the PTB-XL dataset [5].

a) PTB-XL:: According to [5], PTB-XL is the largest open-source ECG dataset since 2020. The dataset contains 12-lead ECG-waveforms from 21837 records of 18885 patients. Each waveform from PTB-XL has a duration of 10 seconds. Two sampling rates are used to collect the data: 100 Hz and 500 Hz. In our experiment, we employ the 100 Hz waveforms. Each 12-lead ECG waveform is associated with one or several classes out of five classes: normal (NORM), conduction disturbance (CD), myocardial infarction (MI), hypertrophy (HYP), and ST/T change (STTC). For waveforms that belong to multiple classes, we choose only the first one and remove the others for simplicity. Figure 2 shows an example of a 12-lead normal heartbeat from the PTB-XL dataset. The dataset is then split into the train-test splits with a ratio of 90%-10%. In summary, after processing, we have a training split of size $[19267, 12, 1000]$ with 19,267 ECG waveform samples, of 12 channels (or leads) and 1,000 timesteps each. The test split's size is $[2163, 12, 1000]$

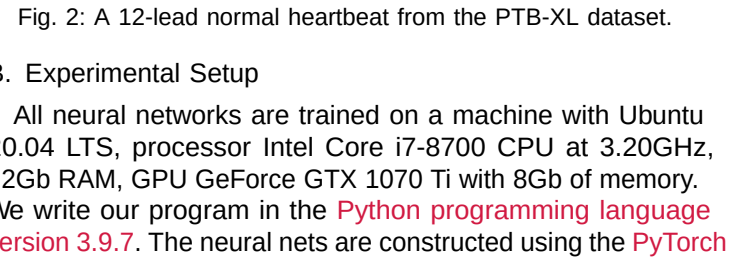

Fig. 2: A 12-lead normal heartbeat from the PTB-XL dataset.

B. Experimental Setup

All neural networks are trained on a machine with Ubuntu 20.04 LTS, processor Intel Core i7-8700 CPU at 3.20GHz, 32Gb RAM, GPU GeForce GTX 1070 Ti with 8Gb of memory. We write our program in the Python programming language version 3.9.7. The neural nets are constructed using the PyTorch library version 1.8.1+cu102. For HE algorithms, we employ the TenSeal library version 0.3.10.

In terms of hyperparameters, we train all networks with 10 epochs, a learning rate of $\eta = 0.001$, and a batch size of $b = 4$. For the split neural network with HE activation maps, we use the Adam optimizer for the client model and mini-batch Gradient Descent for the server model. We use GPU for networks trained on the plaintext. For the U-shaped SL model on HE activation maps, we train the client model on GPU, and the server model on CPU.

C. Evaluation

In this section, we report the experimental results in terms of accuracy, training duration and communication throughput.

The 1D CNN models used for the PTB-XL datasets have two Conv1D layers and one linear layer. The activation maps are the output of the last Conv1D layer.

For the PTB-XL dataset[1], the number of the input channels for the first Conv1D layer are 12, since the input data are 12-lead ECG signals. Besides, we only experiment with 8 output channels for the second Conv1D layer. We denote this

[1]For the MIT-BIH dataset, please have a look at [14]

Fig. 3: Training locally on the plaintext PTB-XL dataset.

network by M. Using M, the output activation map size is [batch size, 2000].

Training the model M locally (without any split) on the plaintext PTB-XL dataset results in Figure 3. As we notice the best training accuracy after 10 epochs is only 72.65%. The best test accuracy is 67.68%. The reason for this low accuracy is that our neural network is small (it contains only 12013 trainable parameters), and we only train it for a particularly small epoch. Each training epoch takes 10.56 seconds on average.

D. U-shaped Split Learning using Plaintext Activation Maps

Based on our experiments, training U-shaped split model on plaintext yields same results in terms of accuracy compared to local training [3]. Although the authors of [3] only used the vanilla version of the split model, they too found that, compared to training locally, accuracy was not reduced.

We will now discuss the training time and communication overhead of the U-shaped split models and compare them to their local versions. For M, the communication overhead is on average 316.9 Mb per epoch, which is much bigger due to the bigger activation maps sent from the client during training.

E. Visual Invertibility

We visualize the activation maps produced by the model M to access their visual similarity compared to the original signals. In SL, the activation maps are sent from client to server to continue the training process. A visual representation of the activation maps reveals a high similarity between certain activation maps and the input data from the client, as shown in Figure 4, Figure 5, Figure 6, Figure 7, Figure 8 for five different classes of heartbeat in the PTB-XL dataset.

The figure indicates that, compared to the raw input data from the client, some activation maps have exceedingly similar patterns. This phenomenon clearly compromises the privacy of the client's raw data. The authors of [3] quantify the privacy leakage by measuring the correlations between the activation maps and the raw input signal by using two metrics: distance correlation and Dynamic Time Warping. This approach allows them to measure whether their solutions mitigate privacy leakage work. Since our work uses HE, said metrics are unnecessary as the activation maps are encrypted.

F. U-shaped Split 1D CNN with Encrypted Activation Maps

The results of training different settings of M on the PTB-XL dataset are in Table I. The HE set of parameters $P = 8192$, $C = [40, 21, 21, 40]$, $\Delta = 2^{21}$ achieves the best test accuracy at 65.42%. This result is only 2.26% lower than the result obtained by the plaintext version. However, this set of parameters incurs the most communication overhead (115.4 Tb per epoch) and takes the longest to train (72 534 seconds per epoch). Overall, test accuracies achieved by different HE parameters are quite close to each other, with 65.42% being the lowest. Interestingly, the smallest set of HE parameters $P = 2048$, $C = [18, 18, 18]$, $\Delta = 2^{16}$ achieves the second-best accuracy at 65.33% while only requiring about 1/10 of the training duration and 1/100 of the communication overhead compared to $P = 8192$. The two sets of parameters with $P = 4096$ produce quite similar results, while roughly taking same amount of time and communication overhead to train.

Fig. 4: Visual invertibility of the model M on the PTB-XL dataset. Left: input data (NORM class). Right: corresponding activation maps.

Fig. 5: Visual invertibility of the model M on the PTB-XL dataset. Left: input data (MI class). Right: corresponding activation maps.

Through our experiments, we see that training on encrypted activation maps can produce very optimistic results, with accuracy dropping by 2-3% for the best sets of HE parameters. Furthermore, training using BE can significantly reduce the amount of training time and communication overhead needed, while producing comparable results when it come to training without BE. The set of parameters with $P = 8192$ always achieve the highest test accuracy, though incurring the highest communication overhead and the longest training time. The set of parameters with $P = 4096$ can offer a good trade-off as they can produce on-par accuracy with $P = 8192$, while requiring significantly less communication and training time. Experimental results show that with the smallest set of HE parameters $P = 2048$, $C = [18, 18, 18]$, $\Delta = 2^{16}$, the least amount of communication and training time is required. In addition, this only works well when used together with BE. When training the network M on the PTB-XL dataset, this set of parameters produces even better test accuracy compared to $P = 4096$. However, this result may be because the network M is small. The test accuracy on the plaintext version is 67.68%, hence the noises produced by the HE algorithm do not yet