

Personalized Real-Time Federated Learning for Epileptic Seizure Detection

Saleh Baghersalimi *IEEE Member*, Tomas Teijeiro *IEEE Member*, David Atienza, *IEEE Fellow*,
Amir Aminifar, *IEEE Member*

Abstract—Epilepsy is one of the most prevalent paroxysmic neurological disorders. It is characterized by the occurrence of spontaneous seizures. About 1 out of 3 patients have drug-resistant epilepsy, thus their seizures cannot be controlled by medication. Automatic detection of epileptic seizures can substantially improve the patient’s quality of life. To achieve a high-quality model, we have to collect data from various patients in a central server. However, sending the patient’s raw data to this central server puts patient privacy at risk and consumes a significant amount of energy. To address these challenges, in this work, we have designed and evaluated a standard federated learning framework in the context of epileptic seizure detection using a deep learning-based approach, which operates across a cluster of machines. We evaluated the accuracy and performance of our proposed approach on the NVIDIA Jetson Nano Developer Kit based on the EPILEPSIAE database, which is one of the largest public epilepsy datasets for seizure detection. Our proposed framework achieved a sensitivity of 81.25%, a specificity of 82.00%, and a geometric mean of 81.62%. It can be implemented on embedded platforms that complete the entire training process in 1.86 hours using 344.34 mAh energy on a single battery charge. We also studied a personalized variant of the federated learning, where each machine is responsible for training a deep neural network (DNN) to learn the discriminative electrocardiography (ECG) features of the epileptic seizures of the specific person monitored based on its local data. In this context, the DNN benefitted from a well-trained model without sharing the patient’s raw data with a server or a central cloud repository. We observe in our results that personalized federated learning provides an increase in all the performance metric, with a sensitivity of 90.24%, a specificity of 91.58%, and a geometric mean of 90.90%.

Index Terms—Deep learning, Electrocardiogram (ECG), Epilepsy, Federated learning (FL), Low-power, Privacy-preserving, Seizure detection

I. INTRODUCTION

Epilepsy is one of the most common neurological disorders affecting around 65 million people worldwide [1]. Among this population, one-third of the patients live with uncontrollable

Manuscript received July 9, 2020; revised August 31, 2020, January 3 2021, and May 17, 2021; accepted July 3, 2021. This work has been partially supported by the ML-Edge Swiss National Science Foundation (NSF) research project (GA No. 200020182009/1), the MyPreHealth Cyber-Human project (Hasler Foundation project No. 16073), the PEDESITE Swiss NSF Sinergia project (GA No. CRSII5_193813 / 1), the H2020 DeepHealth Project (GA No. 825111), the RESoRT project (GA No. REG-19-019) from the Botnar Foundation, and the WASP Program funded by the Knut and Alice Wallenberg Foundation.

Saleh Baghersalimi, Tomas Teijeiro and David Atienza are with the Institute of Electrical Engineering, Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland, (email: saleh.baghersalimi@epfl.ch; tomas.teijeiro@epfl.ch; david.atienza@epfl.ch).

Amir Aminifar is with the Department of Electrical and Information Technology at Lund University, Sweden, (email: amir.aminifar@eit.lth.se)

seizures because of no available treatment working for their cases. One of the most life-threatening effects of epileptic seizures is sudden unexpected death in epilepsy (SUDEP) due to uncontrolled or poorly controlled attacks [2]. Although the percentage of the population affected by SUDEP is relatively low, every death due to SUDEP is thought to be potentially avoidable by notifying caregivers, and emergency units in case of a seizure [3].

To improve the quality of life and control fatal events, patients should be monitored continuously on a long-term basis. This way, in case the patient presents seizures, the medical expert or the family members can be notified for rescue [4, 5]. This long-term every-day remote monitoring is only feasible by using energy-efficient embedded platforms.

Recent development of deep learning has provided a new avenue for addressing these limitations. Deep neural networks (DNNs) have become the new state-of-the-art in the classification of high dimensional data in biomedical applications and recently received increasing attraction in pathology detection. The ability of DNNs to extract high-level and complex abstractions from electrocardiogram (ECG) signals makes them an attractive tool for epilepsy monitoring. However, in order to create a high-quality epileptic seizure detection system, the DNN requires a large-scale dataset. Therefore, training robust DNN classifiers on a single central server often needs the medical data from multiple patients. However, sharing the patient’s raw data with a central server consumes a major amount of energy and, most importantly, puts the patient’s privacy in danger. Privacy law will apply to information if there is a possibility of identifying a specific individual. Moreover, the General Data Protection Regulation (GDPR) [6] is a regulation in the European Union (EU) law on data protection and privacy that addresses the transfer of personal data data outside the EU. Also, the California Consumer Privacy Act (CCPA) is a state law dedicated to enhancing privacy rights and consumer protections for California residents, United States.

In this previous context, Federated Learning (FL) is an excellent option to consider for privacy-related issues in performing data analytics by using multi-institutional big data. It protects patients’ privacy in training of the DNNs by leaving all data within the originating institutions [7, 8]. In FL, a model is learned by multiple clients in a distributed fashion [9, 10]. In the specific case of deep learning, the parameters of the trained network are gathered by a server, and then it distributes an aggregated model back to the clients. The information uploaded by each client is inspectable, but it cannot be used to recover any useful information about the local dataset [11].

In this work, we propose a personalized real-time FL framework in the context of seizure detection on mobile platforms and explore several DNN convolutional neural network architectures to characterize the trade-off between the detection performance and energy consumption. In this proposed framework, we allow every patient to train a personalized DNN by effectively collaborating with the other patients and, while reducing the energy of the communication between the local battery-powered mobile platform and the server. Our proposed approach performs seizure detection based only on the ECG signal from the EPILEPSIAE database and achieves an average sensitivity of 90.24%, specificity of 91.58% and geometric mean of 90.90%. We implement and evaluate our proposed real-time FL framework across an experimental setup with a set of distributed mobile platforms, by including multiple NVIDIA Jetson Nano units [12]. We study and compare the epileptic seizure detection performance in terms of specificity, sensitivity, and geometric mean when the DNNs are trained using centralized approach and the FL framework. Furthermore, we evaluate the energy consumption of the proposed framework. Our experimental results show that local processing (training) and synchronization (communication) can be performed in 1.86 hours using a 344.34 mAh energy on a single charge on NVIDIA Jetson Nano. To the best of our knowledge, this article is the first to study the personalized real-time FL for epileptic seizure detection and by using only ECG signals. The main contributions of our work are summarized below:

- A personalized real-time FL framework in the context of seizure detection on mobile platforms and exploration of DNN convolutional neural network architectures to characterize the trade-off between the detection performance and energy consumption.
- Implementation and demonstration of our proposed real-time FL framework across an experimental setup with a set of distributed mobile platforms in the context of epileptic seizure detection.

The rest of this article is organized as follows. In Section II, we review the latest epileptic seizure detection techniques. In Section III, we describe our personalized real-time FL approach for epileptic seizure detection by defining the elements of our framework and proposed FL technique. Then, in Section IV, the experimental setup is discussed. Next, in Section V, we evaluate the results of our proposed framework, in terms of both detection accuracy and energy consumption on an embedded platform. Finally, in Section VI, we summarize the main conclusions of this work.

II. RELATED WORK

The gold standard in terms of non-invasive seizure detection is EEG monitoring [13], which has been used for decades in highly specialized and costly hospital environments. Several techniques have been developed to extract information from the EEG signal. These techniques use features such as wavelet transform [14, 15], entropies [16], Hilbert marginal

spectrum [17], Hilbert Huang Transform [18], fusion features [19], and tunable Q-factor [20], among other time-domain and frequency-domain features. In addition, different time–frequency based detection methods for epileptic seizures are based on the fact that EEG signals are non-stationary, have been developed [21, 22]. Recently, deep learning has received notable attention in several applications like epileptic focus localization, epileptic seizure prediction and detection. In [23], the authors propose a method for epileptic focus localization problem, where an unsupervised learning scheme is implemented by merging a deep convolutional variational autoencoder and K-means algorithm. In [24], the authors present a pseudo-prospective seizure prediction from EEG signal. They train a deep learning classifier to distinguish between pre-ictal and interictal signals. They employ a convolutional neural network (CNN) to extract significant spatial features from different scalp positions and a recurrent neural network in expecting the incidence of epileptic seizures earlier. In [25], the authors used deep neural networks for epileptic seizures detection using EEG signals with an adaptive implementation of CNNs. In [26], the authors presented a unified multi-view deep learning framework to capture brain abnormalities associated with seizures based on multi-channel scalp EEG signals. In [27], the authors proposed a hybrid bilinear DNN using surface electroencephalogram (sEEG) with an application in the clinical procedures of epilepsy classification diagnosis. Hybrid bilinear models are based on two type of feature extractors namely CNNs and Recurrent Neural Networks (RNNs). CNNs extract spatio-temporal patterns, while RNNs focus on the characteristics of temporal dynamics in input data. However, none of these studies are performed on existing embedded devices, taking into consideration the stringent energy and memory constraints of these devices, and the necessity to permanently wear a cap and the high susceptibility of EEG recordings to artifacts as soon as patients move has not yet enabled the use of this previously mentioned framework for real-time ambulatory chronic monitoring. Furthermore, wearing a cap to monitor EEG outside the hospital is a clear cause of social stigma and discomfort for patients [3].

In addition to abnormal brain activities during epileptic seizures, there is also an effect on the central nervous system manifested on non-EEG signals. In particular, it is well known that ECG abnormalities often occur during or after epileptic seizures and may contribute to SUDEP [28, 29, 30]. Also, several works have shown that in most cases, epileptic seizures are associated with an increase in heart rate [31, 32]. Thus, ECG-based detection of epileptic seizures has recently attracted increasing attention. In [33], the authors proposed a curve-fitting algorithm to characterize the heart rate pattern, as well as a new approach based on a moving median filter to automatic detection of the heart rate changes. Their results lead to a sensitivity of above 90% and a positive predictive value of above 50%, but with a trade-off between the sensitivity and positive predictive values. In [34], the authors proposed a method that combines time and frequency-domain features of heart-rate variability, but was limited to the specific case of

newborns, and reached a sensitivity of 60% and a specificity of 60%. Using a similar approach, in [35], the authors could reach a sensitivity of 70% and a corresponding false-alarm rate of 2.11 per hour. However, none of these studies attempted to find a personalized solution for each user and were performed on an embedded platform, i.e., taking into account the energy consumption of these devices as we consider in this work.

In recent years, there has been a rising interest in developing algorithms to detect epileptic seizures on embedded systems using the ECG signal. In particular, in [36, 37], the authors proposed a multi-parametric machine-learning technique to detect epileptic seizures by analyzing the cardiac and respiratory responses to seizures in the ECG signal. However, their algorithm considered windows of 60 seconds, which increases the latency of real-time seizure detection. In addition, segmenting the dataset into segments of 60 seconds results in having very few samples, leading to insufficient data to train deep learning models.

A significant number of efforts have been dedicated to implement FL algorithms to perform efficient machine learning models. Several studies have conducted training of deep neural networks using FL for various applications, including the scalability of distributed systems and developing methods for learning compressed federated networks from pretrained local network [38, 39, 40, 41]. FL training by iteratively averaging locally trained models has been studied in [42] for DNNs for speech recognition. These works only consider a scenario based on data centers, and do not consider unbalanced data that is fundamental to the FL setting. In the real-world scenario, the local data distribution varies among all patients, and class imbalance problem could result in a slow converge rate of the global model. Therefore, following this line of research, we propose for the first time in this work a real-time FL framework designed for battery-powered mobile platforms for the detection of epileptic seizures. Moreover, we adopt this style of algorithm and perform the appropriate personalized FL, which capture their individual characteristics. Thus, FL researches can be divided into three categories, i.e., horizontal FL proposed by Google [43, 44], vertical FL and federated transfer learning, which are based on the distribution characteristics of the data. In this case, we focus on horizontal FL where datasets in all patients share the same feature space but different sample and allows patients to collaboratively learn a DNN model. At the same time, the proposed framework enables data analysis among hospitals/patient homes without sharing sensitive raw data, which is important for the preservation of patient’s privacy.

III. FEDERATED LEARNING FRAMEWORK

In this section, we describe our proposed personalized real-time FL framework for epileptic seizure detection. As shown in Fig. 1, the pipeline of our FL framework is divided into three parts: ECG signal pre-processing, epileptic seizure detection DNN, and personalized federated learning.



Fig. 1: Pipeline of the proposed FL process for epileptic seizure detection.

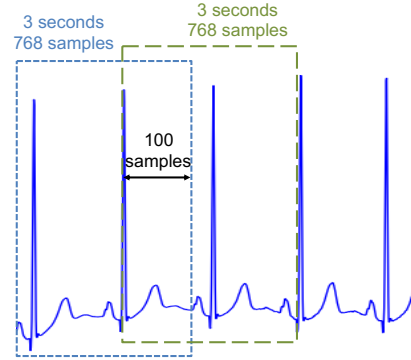


Fig. 2: ECG signal segmentation using slots of 3-seconds with 100 samples overlap.

A. ECG Signal Pre-Processing

A DNN requires a considerable amount of data for the training process. Unfortunately, the amount of data of actual epileptic seizures for a particular patient is generally limited. Thus, we introduced data augmentation by segmenting the ECG signals, which were acquired with a sampling frequency of 256 Hz, into slots of 3-seconds (768 samples). These slots are obtained by sliding a fixed-length window, with 100 samples overlapping, through the entire signal. One of the most important parts of ECG signal processing is the interpretation of QRS complex and obtaining its characteristics. In this work, to ensure to have this complex, we considered 3-seconds windows. Figure 2 shows how the segmentation of ECG signal is performed in our approach.

After ECG segmentation, we propose a simple method for each ECG signal segment’s pre-processing, as shown in Fig. 3. It consists of three steps. First, we apply a 10th-order low-pass Butterworth filter with a cutoff frequency of 50 Hz to smooth the signal segment. Secondly, we perform a linear detrending, where the result of a linear least-squares fit to data is subtracted from the initial data. Finally, we apply standardization on each segment to transform it to have zero mean and unit variance.

B. Epileptic Seizure Detection DNN

In this section, we describe our proposed DNN, referred to as Residual 1-Dimensional Convolutional Neural Network (Res1DCNN). This framework learns to distinguish the seizure from the non-seizure data samples. Given a data sample, the DNN returns an output of size 2 that predicts the label of the data sample, i.e., seizure or non-seizure. We design our

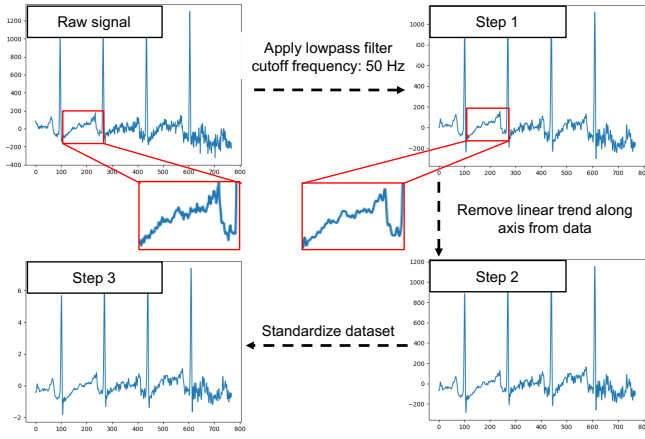


Fig. 3: Pre-processing consists of three steps: (1) applying a low-pass filter, (2) removing linear trend along the axis from data, and (3) standardization of datasets, to have zero mean and unit variance.

proposed DNN model in three steps: starting with Multilayer Perceptron (MLP), refining the model to a 1-Dimensional Convolutional Neural Network (1DCNN), and finally introducing the residual connections to obtain the Residual 1-Dimensional Convolutional Neural Network (Res1DCNN). The final model, i.e., Res1DCNN, is designed to be efficient, in terms of energy consumption and detection performance, for real-time FL on mobile platforms with tight energy operation constraints (cf. Section V).

1) *Multilayer Perceptron (MLP)*: As shown in Fig. 4a, we have developed a multilayer perceptron [45] which is a modified version of an autoencoder proposed by [46] to apply supervised learning. This network takes as input a time-series of pre-processed ECG signals, and they go through 4 hidden layers, which are the computational engine of the network, and outputs a sequence of label predictions. Table I presents the architectural details of this network. The 3-second ECG signal is sampled at 256 Hz. Here, we consider a supervised learning setting, and the network is trained on a set of input and output (label) pairs, where each input is associated with a label defining its class. The main problem in MLP, which is a fully connected network, is that each neuron in one layer is connected to all neurons in the next layer. This fully-connectedness makes the training slower and can cause overfitting. To address this problem, we use a CNN design combined with a fully-connected layer. Our CNN is the results of a linear operation using a subset of the weights of a fully-connected layer. The weights for the convolutions at each location are shared. Then, thanks to the weight sharing and using a subset of weights of a fully connected, fewer parameters are needed in comparison to a fully-connected network.

2) *1-Dimensional Convolutional Neural Network (1DCNN)*: Convolutional neural networks [47] are well known for performing feature extraction and finding simple

TABLE I: Architectural details of the MLP.

Input: 3 seconds signal \equiv (None, 768)		
Layer type	Dimensions	Outputs
Fully connected	768×384	(None, 384)
Fully connected	384×768	(None, 768)
Fully connected	768×384	(None, 384)
Fully connected	384×2	(None, 2)
Output: (None, 2)		

and complicated patterns in the data. The 2D convolutional neural networks are widely used for object detection [48], image classification [49] and semantic image segmentation [50]. Since our data are a one-dimensional signal, we have used 1D-CNN architecture, which is able to accurately extract features from our fixed-length ECG segments, to perform epileptic seizure detection. Figure 4b shows the design of this network, in which four layers of 1D-CNN are used, followed by a fully-connected layer. Table II presents the architectural details of this network.

TABLE II: Architectural details of the 1DCNN network.

Input: 3-seconds signal \equiv (None, 768, 1)		
Layer type	Dimensions	Outputs
Convolutional 1d	10, 1 \rightarrow 10, /2	(None, 380, 10)
Convolutional 1d	7, 10 \rightarrow 100, /2	(None, 187, 100)
Max pooling 1d	2, /1	(None, 93, 100)
Convolutional 1d	3, 100 \rightarrow 200, /2	(None, 46, 200)
Convolutional 1d	3, 200 \rightarrow 200, /2	(None, 22, 200)
Global average pooling 1d		(None, 200)
Fully connected	200×2	(None, 2)
Output: (None, 2)		

3) *Residual 1-Dimensional Convolutional Neural Network (Res1DCNN)*: The architecture of the proposed end-to-end Res1DCNN model is shown in Fig. 4c. It starts with a block, which performs a convolution on the input signal with a kernel size of 7 and a stride of 2, followed by a max pooling with a stride of 2. Then, the output is processed by 4 residual blocks with 2 convolutional layers per block, all convolutional layers have a filter length of 3 and have $64k$ filters, where k starts out as 1 and is incremented in every residual block. In the end, a fully connected layer is included, so that its output is passed to a softmax layer to predict the class of the input signal. Our final solution includes multiple skip connections similar to those found in Residual Neural Networks [51]. These skip connections make residual blocks norm-preserving, thus allowing us to propagate well the information in very deep neural networks to render the training stable. In total, the resulting network consists of 14 weight layers. At the same time, we can implement the proposed network and perform the training process on a mobile platform with limited computational resources and powered by a battery, as described in the next subsection. Table III presents the architectural details of our proposed network.

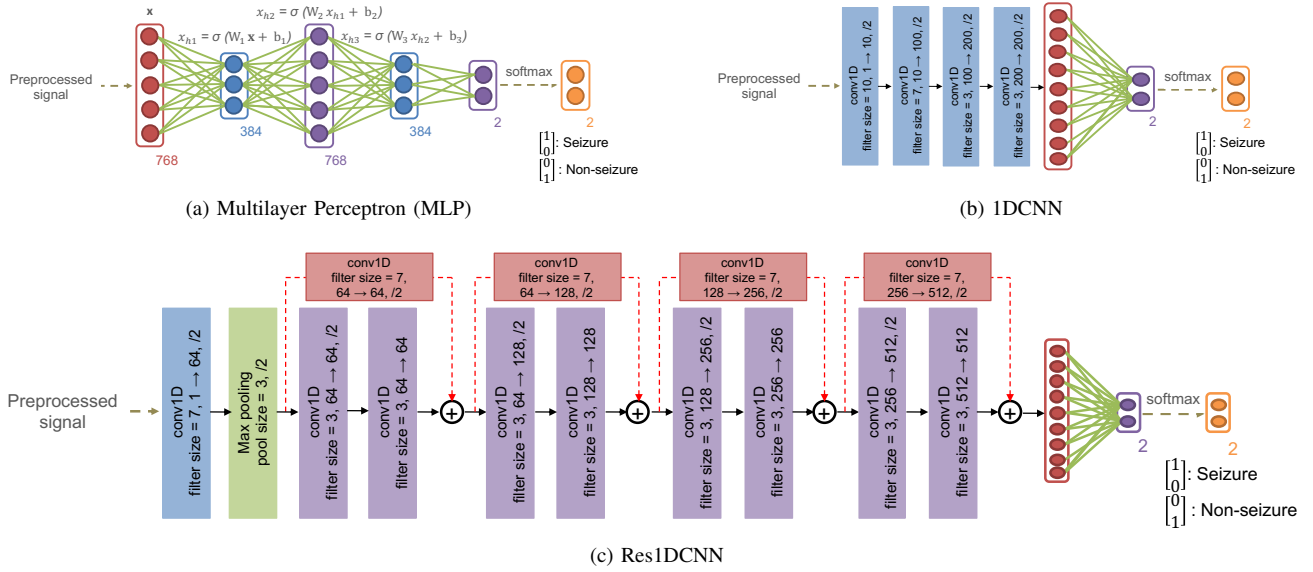


Fig. 4: Our explored networks architecture for epileptic seizure detection. a) The architecture of the Multilayer Perceptron (MLP). This is equivalent to having a 4-layer neural network taking the preprocessed ECG signal. b) The architecture of the 1DCNN. The network contains 4 convolutional layers followed by a fully connected layer and a softmax function. c) The architecture of the proposed Res1DCNN. The network contains 13 convolutional layers with skip connections followed by a fully connected layer and a softmax function.

TABLE III: Architectural details of the proposed Res1DCNN for epileptic seizure detection in mobile platforms.

Input: 3-seconds signal \equiv (None, 768, 1)		
Layer type	Dimensions	Outputs
Convolutional 1d	7, 1 \rightarrow 64, /2	(None, 381, 64)
Max pooling 1d	3, /2	(None, 190, 64)
Convolutional 1d	3, 64 \rightarrow 64, /2	(None, 94, 64)
Convolutional 1d	3, 64 \rightarrow 64, /1	(None, 92, 64)
Convolutional 1d	3, 64 \rightarrow 128, /2	(None, 45, 128)
Convolutional 1d	3, 128 \rightarrow 128, /1	(None, 43, 128)
Convolutional 1d	3, 128 \rightarrow 256, /2	(None, 21, 256)
Convolutional 1d	3, 256 \rightarrow 256, /1	(None, 19, 256)
Convolutional 1d	3, 256 \rightarrow 512, /2	(None, 9, 512)
Convolutional 1d	3, 512 \rightarrow 512, /1	(None, 7, 512)
Fully connected	7 \times 512 \rightarrow 2	(None, 2)
Output: (None, 2)		

C. Personalized Federated Learning

To train DNNs, a large amount of data are required to fit the parameters (often in the order of millions). However, transferring the raw data to a server consumes a significant amount of energy [52, 53] and puts the patient's privacy at risk for medical applications. A promising solution to address this problem is to distribute the computation across several clients (machines), which is known as federated learning (FL) when a central hub coordinates the learning process [38, 39].

In the standard FL framework, the goal is to obtain a global DNN model for all clients. By assuming that all clients' data come from the same (or similar) distribution, we expect that the global DNN results in a better accuracy on any client than the client's own local DNN. However, when the clients' data

come from different distributions and do not necessarily follow the same profile, the global DNN is not adapted for each client.

To consider the profile of each individual patient, while still exploiting the data from other patients, we propose a personalized variant of the FL framework. Figure 5 illustrates our proposed personalized FL framework with n clients in the context of epileptic seizures monitoring. In our proposed FL framework, each client trains their local models, while contributing to the global model maintained by the central server. Each client can take advantage of the information from the global DNN model to compensate for its small training data. When the local distribution is highly correlated with global distribution, the global DNN is preferable; otherwise, the global DNN could be ineffective and the client keeps the local DNN. In Section V-A, we demonstrate that the mixture of local and global DNNs increases the detection performance.

In our personalized FL framework, the training is distributed among the clients, which use the same DNN architecture. In each iteration, the clients are provided with the global model W_i . Then, each client decides whether to exploit the global model received from the server or to perform the training based on its data locally and independently. After the clients are trained for a fixed number of iterations, the server collects the current weights of each client ($W_{i,1}, W_{i,2}, \dots, W_{i,n}$). In the central server, the weights are averaged, and all the clients receive a copy of the updated weights W_{i+1} and continue the training process. Sending the weights to the server, instead of the raw data, significantly reduces the amount of energy consumption in communication between the clients and the server, which is important for battery-powered

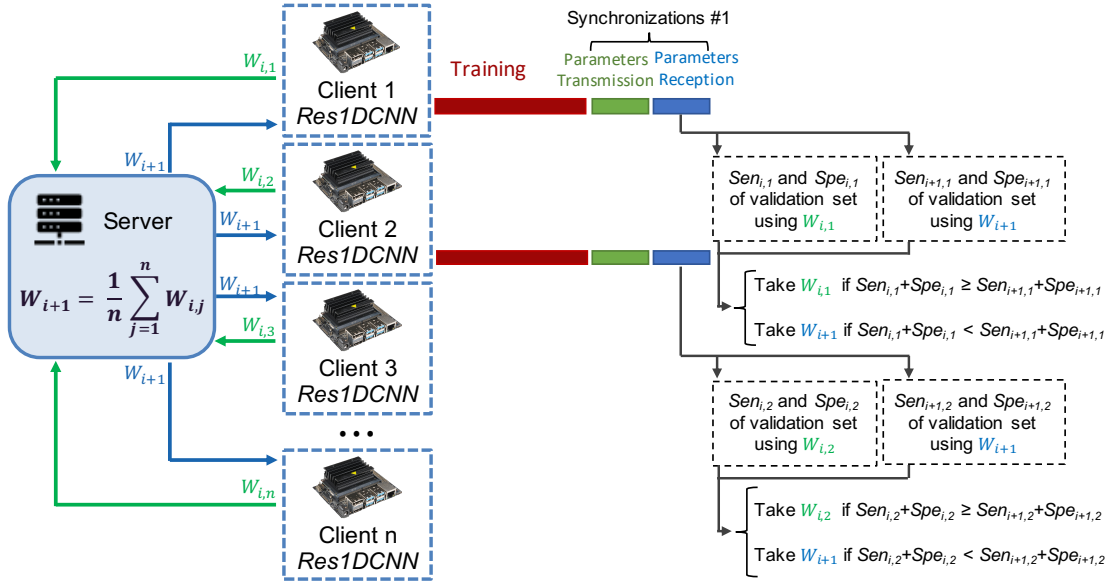


Fig. 5: The overview of our personalized FL system with n clients. Each client trains an identical DNN on its data locally and independently. During the training process, and after a specific number of iterations, each client transmits its current weights to the server. The server calculates the global weights. Then, each client receives the updated weights and locally decides whether to update its personalized weights based on the global weights that it receives from the server.

mobile platforms.

To determine the number of training iterations in our personalized FL framework, we adopt an early stopping technique as a termination criterion. This early stopping technique specifies a maximal number of epochs and retains the model once the model performance stops improving on the validation dataset. Once this criterion is met, we terminate the training process of the clients. In the end, we obtain a model that is personalized for each individual patient, but also robust since it exploits the data from various patients.

Our personalized FL framework also contributes to preserving the privacy of the clients/patients involved in the training process. This is because the sensitive raw personal/medical data that leave the patients' devices are limited in our proposed framework. In particular, the server does not receive any information regarding the DNN model architecture and the clients do not share any information about the number, type, and connectedness of network layers, or the connection type (fully connected layer, convolution layer, deconvolution Layer, pooling layer, etc.) and hyperparameters (stride size and activation functions). As a result, during the FL training process, the server obtains the weights in a pre-defined order from all clients and disseminates the updated global weights back to the clients. For instance, the Res1DCNN contains 13 convolutional layers and a fully connected layer, with a total of 2.8 million weights. Therefore, during the FL training process, each client sends its weights in an pre-defined order as an array $(w_1, w_2, \dots, w_{2,800,000})$ to the server.

Finally, we also explore the practicality of our personalized FL framework on battery-powered mobile platforms with limited resources in Section V-B. In fact, one of the main

challenges in this procedure is the fact that the number of synchronizations cannot be defined strictly and may vary in different setups. On the one hand, in case of frequent synchronizations, the energy consumption of the transmission and reception increases. On the other hand, if this synchronization is performed rarely, the local learning weights for each client will diverge, which results in a poor model after averaging. In Section V-B, we explore the trade-offs between the epileptic seizure detection performance of the model and its energy consumption, as a function of the number of synchronizations.

IV. EXPERIMENTAL SETUP

In this section, we present the experimental setup for the evaluation of our proposed real-time FL framework in terms of detection performance and energy consumption.

A. Epileptic Seizures Dataset

For our experiments, we have used the EPILEPSIAE dataset [54], which is one of the largest epilepsy datasets manually annotated by medical experts for seizure detection and prediction worldwide and enables us to rigorously evaluate our proposed methodology. This dataset consists of one-lead ECG and 19-channel EEG data of 30 patients. The recordings are made in a routine clinical environment, so non-seizure activity and artifacts such as head/body movement, chewing, blinking, early stages of sleep, and electrode pops/movement are present. No constraints regarding the types of seizure are imposed. The dataset contains complex partial (CP), simple partial (SP), and secondarily generalized seizures (GS).

For one of the 30 patients in EPILEPSIAE dataset, different recorded biosignals had different length and as a result it was

not possible to synchronize them to label the data reliably. Thus, we have excluded the signals of this patient from our analyzes and used the ECG data of 29 patients with 4603 hours of recordings containing 277 seizures. The data were acquired at a sampling rate of 256 Hz with 16-bit resolution. We segmented seizure and non-seizure events of each patient into overlapping windows of 3-second and fed them into the proposed architecture of the MLP or CNN. Table IV details the number of segments of seizures and non-seizures extracted for each patient.

TABLE IV: Number of non-seizure and seizures segments per patient

Patient	Non-seizures	Seizures	Patient	Non-seizures	Seizures
#1	209,259	322	#16	124,239	124
#2	346,959	274	#17	237,994	99
#3	215,591	229	#18	194,744	516
#4	190,395	174	#19	216,150	148
#5	236,352	157	#20	208,739	188
#6	180,437	157	#21	214,409	291
#7	121,983	160	#22	127,108	272
#8	150,019	140	#23	183,343	251
#9	152,391	87	#24	311,367	188
#10	180,678	174	#25	174,798	320
#11	184,690	259	#26	212,439	150
#12	210,437	138	#27	210,149	685
#13	212,925	273	#28	212,623	1,001
#14	209,101	280	#29	208,991	338
#15	156,801	121			

We split the dataset into training, validation, and test set. The training set contained 6716 and 5,794,311 segments of seizures and non-seizures. To evaluate the detection accuracy of the proposed DNNs and the influence of the standard FL approach in the classification of each individual segment, we considered a balanced scenario where both the validation and test set include 400 segments of seizures and 400 segments of non-seizures each. Thus, we made sure that no overlap exists between the training, validation, and test set segments. Most importantly, the training, validation, and test set segments of a specific patient are from different 1-hour signal recordings. Table V gives a summary of the considered three sets. Since we have unbalanced training data, we perform undersampling. Undersampling means that from the majority class which is non-seizure, we select as many segments as the minority class which is seizure. This selection maintains the probability distribution of the class during the training process.

In addition, since in practice the classification must be made considering the unbalanced scenario on test set in seizure detection, in Section V-A, we used a leave-one-out cross-validation approach. Thus, the DNN is trained on all the patients except for one patient, and the actual epileptic seizure detection is performed for that patient. In this way, we ensure the temporal and subject independence of the test set, hence, guaranteeing the generality of the results.

TABLE V: Dataset summary

Dataset	Seizure	Not seizure
Training set	6,716	5,794,311
Validation set	400	400
Test set	400	400




Fig. 6: The procedure of FL, which contains two major steps of local training and synchronization with the server.

B. Detection Performance Metric

To evaluate the detection performance of our proposed framework, we considered five metrics: the sensitivity (Sen), which represents the percentage of ictal samples that are labeled correctly; the specificity (Spe), which shows the percentage of inter-ictal samples that are labeled correctly; and the geometric mean (Gmean) [55], which reflects both sensitivity and specificity. The geometric mean measures the balance between classification performance on both classes. A low geometric mean indicates poor performance in the classification of the seizure cases, even if the non-seizures cases are correctly classified. Thus, we also evaluate the accuracy (Acc), which represents the proportion of true positive results in the selected population. Finally, we use the F_1 score, which is the harmonic mean of precision and recall. It gives a better measure of the incorrectly classified cases than the accuracy. These metrics are defined as follows:

$$Sen = \frac{TP}{TP + FN},$$

$$Spe = \frac{TN}{FP + TN},$$

$$Gmean = \sqrt{Sensitivity \times Specificity},$$

$$Acc = \frac{TP + TN}{TP + TN + FP + FN},$$

$$F_1score = \frac{2TP}{2TP + FP + FN},$$

where TP, TN, FP and FN are true positive, true negative, false positive and false negative, respectively.

We also use a confusion matrix [56] to evaluate the detection performance of our proposed framework. Since the output can be one of two types of classes, the confusion matrix




Fig. 7: Comparison of the epileptic seizure detection performance of our explored DNNs trained using a centralized approach versus federated learning setting.

is one of the most intuitive and easy metrics used for this type of classification tasks. In this case, the diagonal elements represent the number of cases for which the predicted label is equal to the correct label, while off-diagonal elements are those that are mislabeled by the classifier.

C. Implementation Platform

We use the NVIDIA Jetson Nano Developer Kit GPU [12] for the training. It allows us to perform the training of a given DNN with a power consumption of only 5 watts. The NVIDIA Jetson Nano includes a 64-bit quad-core Arm Cortex-A57 CPU running at 1.43 GHz alongside a NVIDIA Maxwell GPU with 128 CUDA cores capable of 472 GFLOPs (FP16), and has 4 GB of 64-bit LPDDR4 RAM. We used Tegrastats [57] for performance analysis and characterization of training DNNs on NVIDIA Jetson Nano. Tegrastats is a tool provided by NVIDIA that collects hardware utilization, memory and power consumption of both CPU and GPU.

D. Learning Parameters

We trained our proposed networks from scratch using pre-processed 3 seconds ECG segments. The weights initialization of the layers follows a normal distribution with zero mean and 0.01 as standard deviation. We initialize all the biases to zero. During the training, the network learns the correlation between the input and the output consisting of two nodes and adjusts the parameters of the model to minimize the cross-entropy loss. For a binary classification, the final output of DNN can have a single output and a threshold or we can use a multi-class classification with only two nodes, so each class gets its output neuron. However, the two node outputs technique code is exactly the same for multi-class classification problem, and can be easily extended to multi-class classification in future

work. Finally, we use the Adam optimizer with a base learning rate of 10^{-5} and implement the DNNs on Tensorflow 1.14.0 [58].

E. Distributed Training Parameters

As shown in Fig. 6, to evaluate the epileptic seizure detection performance and energy consumption of our proposed DNNs for FL, we considered a random division of the patients data between 4 clients that hold the same model and perform the training process separately. Then, after the local models are trained for 3,000 iterations, the server collects all the models from each client and the weights are averaged. The clients receive a copy of the updated weights and continue the training process. The number of iterations between two consecutive synchronizations has to be selected carefully. In Section V-B, we explore that if we fix the number of iterations to be very low, the energy consumption of communication increases. If we increase the number of iterations to be high, the weights in each client will diverge significantly. In other words, the client's model will have different local minima and averaging the weights will not guarantee to find the global-local minima. Therefore, the accuracy of the final model will be reduced.

To assess the benefits of our proposed personalized FL approach with respect to a standard FL framework within a real-world scenario (i.e., where there are many clients), we also considered a set-up with 29 clients. Each client contains the data of only one patient and holds the same DNN model. In personalized FL, when the clients receive a copy of the updated global weights, they compare the epileptic seizure detection performance on the local validation data using the updated global weights versus using their local weights and can decide to accept the weights, thus resulting in a higher detection performance. Each client allocates 80% of its data




Fig. 8: ROC curve of our explored DNNs trained using a centralized approach versus FL setting. The area under the ROC curve while using the centralized approach and using FL for Res1DCNN are 0.9113 and 0.8956, while for 1DCNN, it decreases from 0.8760 to 0.8378, and from 0.8081 to 0.7851 for MLP.

for training and the remaining 20% data are equally partitioned between validation and test set.

V. EVALUATION

In this section, we present the assessment of the seizure detection performance and energy consumption of our proposed real-time federated-learning framework. In particular, as a case study, we explore a distributed mobile platform setup, including multiple NVIDIA Jetson Nano units [12].

A. Detection Performance Analysis

To verify the effectiveness of our proposed Res1DCNN in terms of epileptic seizure detection performance, we compare it with two other proposed architectures, namely MLP and 1DCNN. Figure 7 presents the detection performance of all the proposed networks when they are trained in a centralized approach without using FL technique. We observe that our proposed 1DCNN achieves the geometric mean of 79.84% and outperforms the MLP by 4%. This shows that 1DCNN is more effective to extract features from the 3-second ECG input signals compared to the MLP architecture. Res1DCNN, which is the combination of 1DCNN and skip connection similar to those found in Residual Neural Networks [51], results in a geometric mean of 84.08% and improvement of 4% compared to 1DCNN. This is due to the fact that these skip connections allow the information to correctly propagate in deeper layers, as we have explored experimentally in this work.

To have an unbiased epileptic seizure detection performance estimation of Res1DCNN, we perform the leave-one-out-cross-validation where the number of folds equals the number

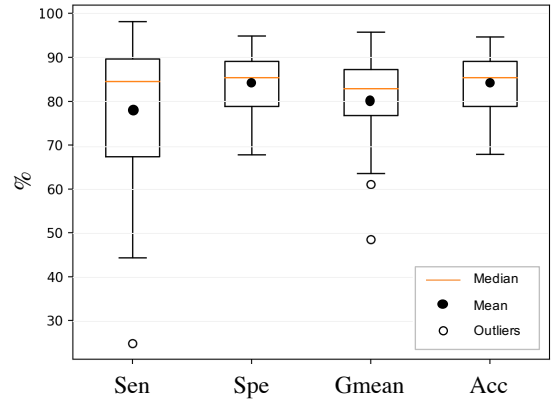


Fig. 9: Box plot of leave-one-patient-out cross-validation of Res1DCNN using only 3 seconds of ECG signal.

of patients, which is 29 in our data set. Thus, the learning algorithm is applied once for each patient, using all other patients as a training set and using the selected patient as a single-item test set. Figure 9 shows the results of the leave-one-out-cross-validation per patient of Res1DCNN while using only 3 seconds of ECG signal. The average sensitivity and specificity are 77.35% and 83.77%, respectively, which results in a geometric mean of 80.49%.

As shown in Fig. 10, we also evaluate the epileptic detection performance when we combine the output of Res1DCNN for multiple consecutive 3 seconds segments of the ECG signal. Figure 11 shows the results of the leave-one-out cross-validation per patient of Res1DCNN while we use multiple consecutive 3 seconds segments, which correspond to 768 samples with an overlap of 100 samples. We observe that by using few more successive segments, we can improve the real-time epileptic seizure detection performance. However, extending the number of segments beyond a certain limit reduces the detection performance, mainly due the fact that the interval under analysis may contain both seizure and non-seizure segments.

Figure 7 shows the results of all the proposed DNNs using the FL setting versus centralized approach. We observe that, in FL setting, the Res1DCNN achieves a geometric mean of 81.62% and outperforms both MLP and 1DCNN by 7% and 6%, respectively.

Then, in Table VI, we compare the normalized confusion matrices of each DNN when they are trained using centralized approach versus the FL setting, and understand the trade-off between the training efficiency and model detection accuracy. The final detection accuracy that we achieve using the FL setting is slightly lower than the detection accuracy when the training process is performed using centralized approach. In particular, in Res1DCNN, we detect 82% of epileptic seizures segments when the training is performed using centralized approach, while in the FL setting, this value decreases by only 1%. Figure 8 compares the ROC curve of the proposed DNNs using the FL setting versus the centralized approach.

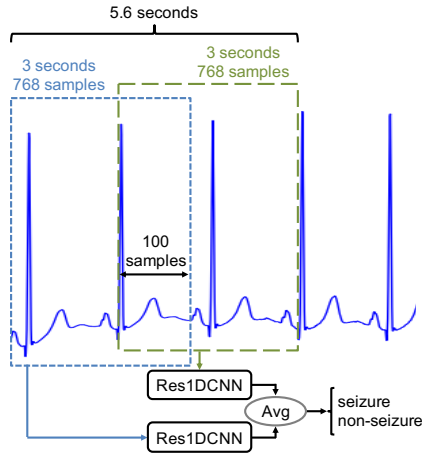


Fig. 10: Example combining the output of Res1DCNN for two consecutive segments of 3 seconds of the ECG signal corresponding to 768 samples with an overlap of 100 samples.

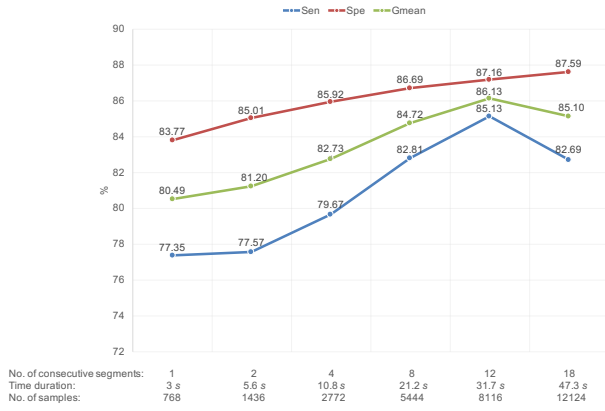


Fig. 11: Average value of sensitivity, specificity and geometric mean, while performing the leave-one-patient-out cross-validation of Res1DCNN, using multiple consecutive 3 seconds segments.

We observe that the area under the ROC curve of Res1DCNN in FL is 0.8956, while in the centralized approach, it is 0.9113. In our view, this situation is the result of the fact that, during the training using centralized approach, we process one batch at a time, then compute the model gradients, and update the weights. On the other hand, when we use our proposed FL setting, we process multiple batches at once, then average the weights of the clients. In the centralized setting, the clients have to trade their privacy by sending the raw data to a central server in order to obtain a slightly better machine learning model, but the FL technique enables the clients at different locations to collaboratively learn a machine learning model, while not exposing their data that may contain private information. In other words, the clients benefit from obtaining a well-trained model without sharing their sensitive personal medical data with other entities, e.g., any server or cloud

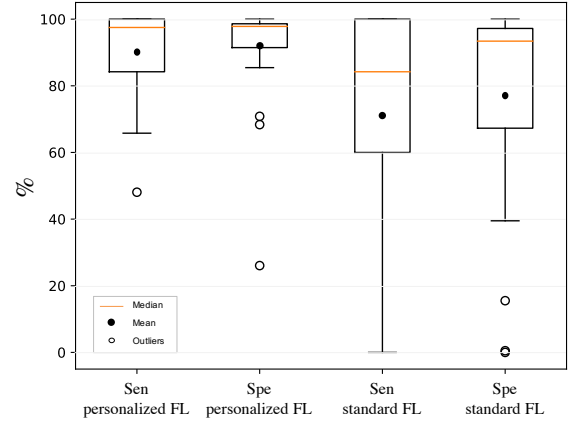


Fig. 12: Box plot of personalized versus standard FL approach using Res1DCNN with 29 clients using only 3 seconds of ECG signal. While using personalized FL, the average sensitivity and specificity are 90.24% and 91.58%, which outperform standard FL by 18% and 14%, respectively.

computing infrastructure.

The results comparing our proposed personalized FL approach with standard FL are depicted in Fig. 12, where we have 29 clients with different amounts of data. Each client contains the data of one patient. We observe that, in the personalized FL setting, the average of sensitivity, specificity, and geometric-mean are 90.24%, 91.58%, and 90.90%. In contrast, in the standard FL setting, these values drop to 72.31%, 77.97%, and 75.08%, respectively.

TABLE VI: Comparison of the normalized confusion matrices of MLP, 1DCNN and Res1DCNN in epileptic seizure detection using centralized approach versus our federated learning setting.

	Centralized approach		Federated Learning	
	Seizure	Non seizure	Seizure	Non seizure
Seizure	79%	21%	76%	24%
Non seizure	27%	73%	28%	72%

MLP

	Centralized approach		Federated Learning	
	Seizure	Non seizure	Seizure	Non seizure
Seizure	78%	22%	69%	31%
Non seizure	18%	82%	18%	82%

1DCNN

	Centralized approach		Federated Learning	
	Seizure	Non seizure	Seizure	Non seizure
Seizure	82%	18%	81%	19%
Non seizure	13%	87%	18%	82%

Res1DCNN

B. Energy Consumption Analysis

The FL scheme relies on communication between the clients and the server to learn a machine learning model. Therefore,

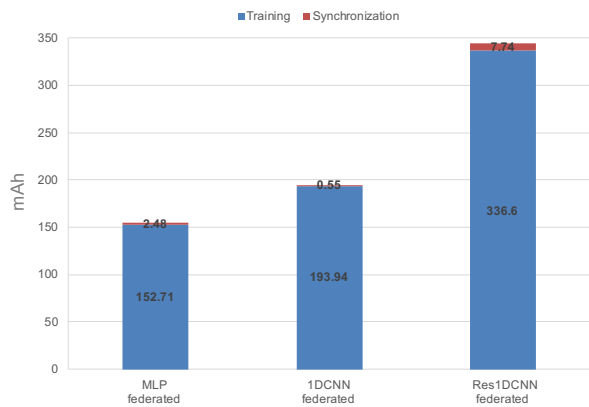


Fig. 13: Required energy capacity of a battery for each client to train the proposed DNN on a single charge in the FL setting, where the server collects the parameters of each client every 3,000 iterations. MLP, 1DCNN and Res1DCNN can be implemented on embedded platforms that complete the entire training process for 1.49, 1.44 and 1.86 hours, respectively.

the first and key challenge in any FL approach is the energy consumption for synchronization of the clients (or machines). In addition, since we would like to implement the epileptic seizure detection algorithm on an embedded medical platform that runs on a battery, the second challenge is the energy consumption for training process. In this context, Fig. 13 shows the energy consumption of each DNN for the training process and synchronization. As this table highlights, we have addressed both challenges by carefully designing our Res1DCNN with only 2.8 million parameters. Hence, we are able to perform the full training of our model on an mobile platform using 344.34 mAh energy for 1.86 hours on a single battery charge.

Let us now consider the Res1DCNN, which outperformed both MLP and 1DCNN, and evaluate its epileptic seizure detection performance versus the number of iterations that the clients apply the synchronization with server to update the weights. We expect the detection performance to increase when the synchronization between the clients and the server is more frequent. On the other hand, the energy consumption of synchronization increases as well, which is not ideal for the embedded platforms running on battery. Indeed, as shown in Fig. 14, if we increase the number of iterations between synchronizations from 30 to 10,000, the total energy consumption referring to both training and synchronization decreases from 1111.40 mAh to 338.98 mAh, but there is also a degradation in terms of geometric mean from 82.36% to 76.74%. Therefore, the number of iterations between two consecutive synchronizations should be carefully selected to be able to use our embedded platform running on battery, while maintaining the epileptic seizure detection performance. As a result, in the context of our experiments, we consider the synchronization every 3,000 iterations, making the implementation possible on mobile battery-powered devices, without any major machine-




Fig. 14: Comparison of the trade-offs between epileptic seizure detection performance and energy consumption of the new Res1DCNN using our proposed FL technique, while the server performs a synchronization every 30, 3000, 5000, 8000, and 10000 iterations.

learning performance loss.

VI. CONCLUSION

In this article, we have designed and implemented a new FL framework for epileptic seizure detection across a cluster of mobile platforms. First, we have proposed an efficient DNN solution that is implementable on low-power embedded medical platform to extract the discriminative ECG features of an epileptic seizure, without sharing the sensitive personal medical data with a server or a cloud. Second, we have developed and validated a new FL approach for epileptic seizure detection using the EPILEPSIAE database on a distributed mobile platform setup, including multiple NVIDIA Jetson Nano units. Our results indicate that this new approach achieves a sensitivity of 81.25% and a specificity of 82.00%. Moreover, our experimental evaluations have proven that the proposed framework can be implemented on a mobile platform powered by a 344.34 mAh battery, thus achieving up to 1.86 hours of operation for a typical battery size used today in mobile medical devices. We have studied a personalized variant of standard FL, where our goal is to find a suitable DNN model for each client that can be adapted to their local data. Our results illustrate that this personalized FL achieves an average sensitivity of 90.24% and specificity of 91.58%. Although the provided analyses and methodologies constitute a set of powerful tools to guarantee the real-time personalized FL framework for epileptic seizure detection, there are some improvements that can be made. In this context, we survey some of the provided results, which can be improved further, as well as we briefly describe some interesting research topics, that are worth investigating further: Using other bio-signal such as EEG combined with ECG signal, and implement the proposed personalized FL framework across other platforms. We are also working on translating the proposed framework

into a real-life scenario by performing ECG and EEG signal acquisition using hardware platforms such as Polar H10 heart rate sensor and e-Glass.

REFERENCES

- [1] Epilepsy. [Online]. Available: URL: http://www.who.int/mental_health/neurology/epilepsy/en/.
- [2] David J Thurman, Dale C Hesdorffer, and Jacqueline A French. "Sudden unexpected death in epilepsy: assessing the public health burden." In: *Epilepsia* 55 10 (2014), pp. 1479–85.
- [3] Anouk Van de Vel et al. "Non-EEG seizure detection systems and potential SUDEP prevention: State of the art Review and update". In: *Seizure* 41 (2016), pp. 141–153.
- [4] Judith van Andel et al. "Non-EEG based ambulatory seizure detection designed for home use: What is available and how will it influence epilepsy care?" In: *Epilepsy Behavior* 57 (2016), pp. 82–89.
- [5] Philippe Ryvlin et al. "Wearable devices for sudden unexpected death in epilepsy prevention." In: *Epilepsia* 59 Suppl 1 (2018), pp. 61–66.
- [6] Jan Philipp Albrecht. "How the GDPR will change the world". In: *Eur. Data Prot. L. Rev.* 2 (2016), p. 287.
- [7] Wenrui Dai et al. "Privacy preserving federated big data analysis". In: *Guide to Big Data Applications*. Springer, 2018, pp. 49–82.
- [8] Qiang Yang et al. "Federated machine learning: Concept and applications". In: *ACM Transactions on Intelligent Systems and Technology (TIST)* 10.2 (2019), pp. 1–19.
- [9] Jie Xu and Fei Wang. "Federated learning for healthcare informatics". In: *arXiv preprint arXiv:1911.06270* (2019).
- [10] Tian Li et al. "Federated learning: Challenges, methods, and future directions". In: *IEEE Signal Processing Magazine* 37.3 (2020), pp. 50–60.
- [11] Qinbin Li et al. "A survey on federated learning systems: vision, hype and reality for data privacy and protection". In: *arXiv preprint arXiv:1907.09693* (2019).
- [12] Jetson Nano Developer Kit: URL: <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>.
- [13] Renée A Shellhaas. "Continuous long-term electroencephalography: the gold standard for neonatal seizure diagnosis". In: *Seminars in Fetal and Neonatal Medicine*. Vol. 20. 3. Elsevier, 2015, pp. 149–153.
- [14] D. Wang et al. "Epileptic Seizure Detection in Long-Term EEG Recordings by Using Wavelet-Based Directed Transfer Function". In: *IEEE Transactions on Biomedical Engineering* 65.11 (Nov. 2018), pp. 2591–2599. ISSN: 1558-2531. DOI: 10.1109/TBME.2018.2809798.
- [15] Ling Xu Guo, Daniel Rivero, and Alejandro Pazos. "Epileptic seizure detection using multiwavelet transform based approximate entropy and artificial neural networks". In: *Journal of Neuroscience Methods* 193 (2010), pp. 156–163.
- [16] S. Raghu et al. "A Novel Approach for Real-Time Recognition of Epileptic Seizures Using Minimum Variance Modified Fuzzy Entropy". In: *IEEE Transactions on Biomedical Engineering* 65.11 (Nov. 2018), pp. 2612–2621. ISSN: 1558-2531. DOI: 10.1109/TBME.2018.2810942.
- [17] Kai Fu et al. "Hilbert marginal spectrum analysis for automatic seizure detection in EEG signals". In: *Biomed. Signal Proc. and Control* 18 (2015), pp. 179–185.
- [18] Yissel Rodriguez Aldana et al. "Nonconvulsive epileptic seizure detection in scalp EEG using multiway data analysis". In: *IEEE journal of biomedical and health informatics* 23.2 (2018), pp. 660–671.
- [19] Jiang-Ling Song, Wenfeng Hu, and Rui Zhang. "Automated detection of epileptic EEGs using a novel fusion feature and extreme learning machine". In: *Neurocomputing* 175 (2016), pp. 383–391.
- [20] Ahnaf Rashik Hassan, Siuly Siuly, and Yanchun Zhang. "Epileptic seizure detection in EEG signals using tunable-Q factor wavelet transform and bootstrap aggregating". In: *Computer methods and programs in biomedicine* 137 (2016), pp. 247–259.
- [21] Alexandros T Tzallas, Markos G Tsipouras, and Dimitrios I Fotiadis. "Automatic seizure detection based on time-frequency analysis and artificial neural networks". In: *Computational Intelligence and Neuroscience* 2007 (2007).
- [22] Alexandros T Tzallas, Markos G Tsipouras, and Dimitrios I Fotiadis. "Epileptic seizure detection in EEGs using time-frequency analysis". In: *IEEE transactions on information technology in biomedicine* 13.5 (2009), pp. 703–710.
- [23] Hisham Daoud and Magdy Bayoumi. "Deep Learning Approach for Epileptic Focus Localization". In: *IEEE Transactions on Biomedical Circuits and Systems* (2019).
- [24] Hisham Daoud and Magdy A Bayoumi. "Efficient epileptic seizure prediction based on deep learning". In: *IEEE transactions on biomedical circuits and systems* 13.5 (2019), pp. 804–813.
- [25] Ömer Türk and Mehmet Siraç Özerdem. "Epilepsy Detection by Using Scalogram Based Convolutional Neural Network from EEG Signals". In: *Brain sciences*. 2019.
- [26] Ye Yuan et al. "A multi-view deep learning framework for EEG seizure detection". In: *IEEE journal of biomedical and health informatics* 23.1 (2018), pp. 83–94.
- [27] Tennison Liu et al. "Epileptic Seizure Classification with Symmetric and Hybrid Bilinear Models". In: *IEEE Journal of Biomedical and Health Informatics* (2020).
- [28] Maromi Nei, Reginald T Ho, and Michael R Sperling. "EKG abnormalities during partial seizures in refractory epilepsy". In: *Epilepsia* 41.5 (2000), pp. 542–548.
- [29] Christian Opherck, James Coromilas, and Lawrence J. Hirsch. "Heart rate and EKG changes in 102 seizures:

- analysis of influencing factors”. In: *Epilepsy Research* 52 (2002), pp. 117–127.
- [30] Marshall J Keilson et al. “ECG abnormalities in patients with epilepsy”. In: *Neurology* 37.10 (1987), pp. 1624–1624.
- [31] LD Blumhardt, PEM Smith, and Lynne Owen. “Electrocardiographic accompaniments of temporal lobe epileptic seizures”. In: *The Lancet* 327.8489 (1986), pp. 1051–1056.
- [32] H Mayer et al. “EKG abnormalities in children and adolescents with symptomatic temporal lobe epilepsy”. In: *Neurology* 63.2 (2004), pp. 324–328.
- [33] Wouter J.C. van Elmpt et al. “A model of heart rate changes to detect seizures in severe epilepsy”. In: *Seizure* 15 (2006), pp. 366–375.
- [34] OM Doyle et al. “Heart rate based automatic seizure detection in the newborn”. In: *Medical engineering & physics* 32.8 (2010), pp. 829–839.
- [35] Kaat Vandecasteele et al. “Automated epileptic seizure detection based on wearable ECG and PPG in a hospital environment”. In: *Sensors* 17.10 (2017), p. 2338.
- [36] Farnaz Forooghifar et al. “A self-aware epilepsy monitoring system for real-time epileptic seizure detection”. In: *Mobile Networks and Applications* (2019), pp. 1–14.
- [37] F. Forooghifar, A. Aminifar, and D. Atienza. “Resource-Aware Distributed Epilepsy Monitoring Using Self-Awareness From Edge to Cloud”. In: *IEEE Transactions on Biomedical Circuits and Systems* (2019), pp. 1–1. ISSN: 1940-9990. DOI: 10.1109/TBCAS.2019.2951222.
- [38] Ruben Mayer and Hans-Arno Jacobsen. “Scalable Deep Learning on Distributed Infrastructures: Challenges, Techniques and Tools”. In: *ArXiv abs/1903.11314* (2019).
- [39] Praneeth Vepakomma et al. “Split learning for health: Distributed deep learning without sharing raw patient data”. In: *CoRR abs/1812.00564* (2018). arXiv: 1812.00564. URL: <http://arxiv.org/abs/1812.00564>.
- [40] Yang Liu et al. “A secure federated transfer learning framework”. In: *IEEE Intelligent Systems* 35.4 (2020), pp. 70–82.
- [41] Mikhail Yurochkin et al. “Bayesian nonparametric federated learning of neural networks”. In: *International Conference on Machine Learning*. PMLR. 2019, pp. 7252–7261.
- [42] Sixin Zhang, Anna Choromanska, and Yann LeCun. “Deep learning with elastic averaging SGD”. In: *arXiv preprint arXiv:1412.6651* (2014).
- [43] H Brendan McMahan et al. “Federated learning of deep networks using model averaging”. In: *arXiv preprint arXiv:1602.05629* (2016).
- [44] Brendan McMahan et al. “Communication-efficient learning of deep networks from decentralized data”. In: *Artificial Intelligence and Statistics*. PMLR. 2017, pp. 1273–1282.
- [45] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. “Multilayer feedforward networks are universal approximators”. In: *Neural networks 2.5* (1989), pp. 359–366.
- [46] Ali Mousavi, Ankit B Patel, and Richard G Baraniuk. “A deep learning approach to structured signal recovery”. In: *2015 53rd annual allerton conference on communication, control, and computing (Allerton)*. IEEE. 2015, pp. 1336–1343.
- [47] Wei Zhang et al. “Shift-invariant pattern recognition neural network and its optical architecture”. In: *Proceedings of annual conference of the Japan Society of Applied Physics*. 1988.
- [48] Joseph Redmon et al. “You only look once: Unified, real-time object detection”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 779–788.
- [49] Kaiming He et al. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2016, pp. 770–778.
- [50] Jonathan Long, Evan Shelhamer, and Trevor Darrell. “Fully convolutional networks for semantic segmentation”. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2015, pp. 3431–3440.
- [51] K. He et al. “Deep Residual Learning for Image Recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016, pp. 770–778. DOI: 10.1109/CVPR.2016.90.
- [52] Hossein Mamaghanian et al. “Compressed sensing for real-time energy-efficient ECG compression on wireless body sensor nodes”. In: *IEEE Transactions on Biomedical Engineering* 58.9 (2011), pp. 2456–2466.
- [53] Farnaz Forooghifar, Amir Aminifar, and David Atienza. “Resource-aware distributed epilepsy monitoring using self-awareness from edge to cloud”. In: *IEEE transactions on biomedical circuits and systems* 13.6 (2019), pp. 1338–1350.
- [54] Matthias Ihle et al. “EPILEPSIAE - A European Epilepsy Database”. In: *Comput. Methods Prog. Biomed.* 106.3 (June 2012), pp. 127–138. ISSN: 0169-2607. DOI: 10.1016/j.cmpb.2010.08.011. URL: <http://dx.doi.org/10.1016/j.cmpb.2010.08.011>.
- [55] Philip J Fleming and John J Wallace. “How not to lie with statistics: the correct way to summarize benchmark results”. In: *Communications of the ACM* 29.3 (1986), pp. 218–221.
- [56] Stephen V Stehman. “Selecting and interpreting measures of thematic classification accuracy”. In: *Remote sensing of Environment* 62.1 (1997), pp. 77–89.
- [57] Nvidia Tegrastats utility: URL: <https://docs.nvidia.com/jetson/14t/index.html>.
- [58] Martin Abadi et al. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: <https://www.tensorflow.org/>.