



## PAPER

## Edge deep learning for neural implants: a case study of seizure detection and prediction

Xilin Liu<sup>1,\*</sup>  and Andrew G Richardson<sup>2</sup><sup>1</sup> Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA, United States of America<sup>2</sup> Department of Neurosurgery, University of Pennsylvania, Philadelphia, PA, United States of America

\* Author to whom any correspondence should be addressed.

E-mail: [xilinliu@seas.upenn.edu](mailto:xilinliu@seas.upenn.edu)**Keywords:** neural interface, machine learning, deep learning, edge computing, CNN, LSTM, seizure detectionRECEIVED  
1 December 2020REVISED  
15 March 2021ACCEPTED FOR PUBLICATION  
1 April 2021PUBLISHED  
26 April 2021**Abstract**

*Objective.* Implanted devices providing real-time neural activity classification and control are increasingly used to treat neurological disorders, such as epilepsy and Parkinson's disease. Classification performance is critical to identifying brain states appropriate for the therapeutic action (e.g. neural stimulation). However, advanced algorithms that have shown promise in offline studies, in particular deep learning (DL) methods, have not been deployed on resource-restrained neural implants. Here, we designed and optimized three DL models for edge deployment and evaluated their inference performance in a case study of seizure detection. *Approach.* A deep neural network (DNN), a convolutional neural network (CNN), and a long short-term memory (LSTM) network were designed and trained with TensorFlow to classify ictal, preictal, and interictal phases from the CHB-MIT scalp EEG database. A sliding window based weighted majority voting algorithm was developed to detect seizure events based on each DL model's classification results. After iterative model compression and coefficient quantization, the algorithms were deployed on a general-purpose, off-the-shelf microcontroller for real-time testing. Inference sensitivity, false positive rate (FPR), execution time, memory size, and power consumption were quantified. *Main results.* For seizure event detection, the sensitivity and FPR for the DNN, CNN, and LSTM models were 87.36%/0.169 h<sup>-1</sup>, 96.70%/0.102 h<sup>-1</sup>, and 97.61%/0.071 h<sup>-1</sup>, respectively. Predicting seizures for early warnings was also feasible. The LSTM model achieved the best overall performance at the expense of the highest power. The DNN model achieved the shortest execution time. The CNN model showed advantages in balanced performance and power with minimum memory requirement. The implemented model compression and quantization achieved a significant saving of power and memory with an accuracy degradation of less than 0.5%. *Significance.* Inference with embedded DL models achieved performance comparable to many prior implementations that had no time or computational resource limitations. Generic microcontrollers can provide the required memory and computational resources, while model designs can be migrated to application-specific integrated circuits for further optimization and power saving. The results suggest that edge DL inference is a feasible option for future neural implants to improve classification performance and therapeutic outcomes.

**1. Introduction**

Many brain injuries and diseases may be treated by implanted devices that provide real-time classification of neural activity to produce a suitable control output. For example, closed-loop neuromodulatory devices control neural activity classified as

pathological with electrical stimulation to treat epilepsy and Parkinson's disease [1]. Brain-machine interface (BMI) devices classify volitional intent to control external communication or movement devices for paralyzed individuals [2]. In both cases, the neural implants record brain activity, select and extract relevant activity features, and perform

classification and control on the basis of these features.

Performance of these devices is largely dependent on the feature selection and classification algorithms. Feature selection aims to transform the often noisy, correlated signals from many recording channels into a few non-redundant, informative inputs to the machine learning classifier. Feature selection is often a manually-specified, time-consuming process that requires substantial domain expertise [3]. Furthermore, current implantable devices such as the NeuroPace RNS and Medtronic Activa PC+S closed-loop neurostimulators have restrained computation resources. Thus, only simple classification algorithms, such as feature thresholding or linear discriminant analysis, have been implemented [4].

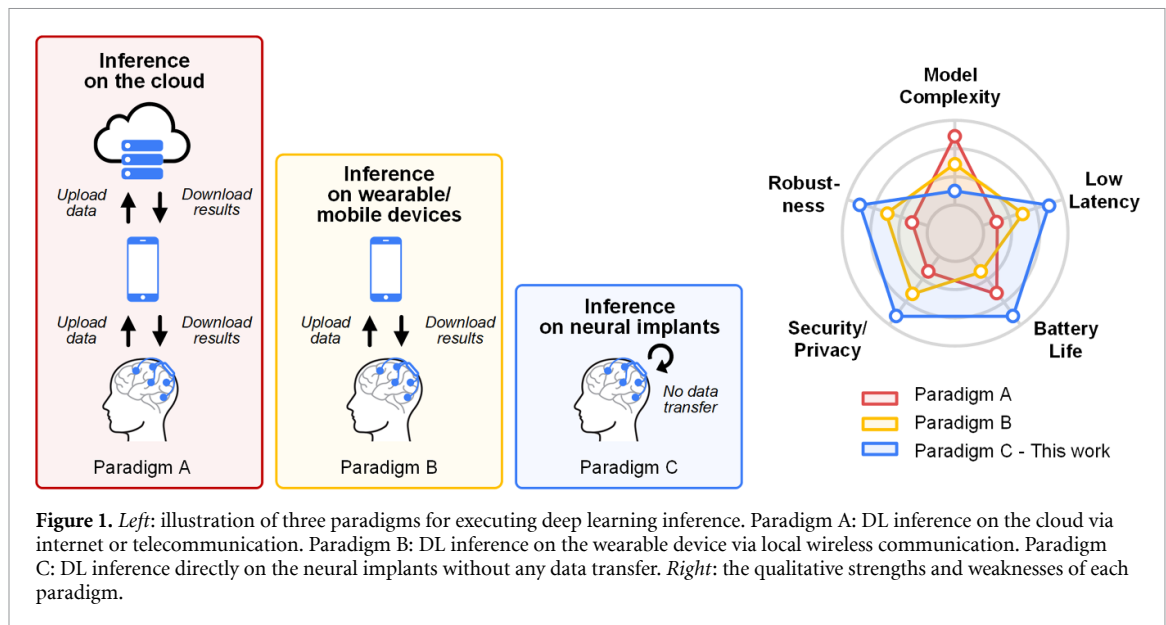
Deep learning (DL) is a machine learning algorithm that has recently been applied to research-grade, non-implantable neural interface devices to improve performance [5]. Specifically, DL combines feature extraction, feature selection, and classification into a single framework, jointly optimizing the end-to-end process [6]. When sufficient training data is available [7], DL can achieve superior performance compared to more conventional algorithms [8], especially in distinguishing hidden features critical for classification [9]. Furthermore, the DL approach is robust and more generalizable across different applications [3]. Although numerous DL approaches for neural interface devices have been studied, nearly all have done so offline [7]. While DL training is computationally intensive and likely to remain offline, DL inference must be performed online for real-time control. This is a particularly challenging problem for clinical-grade devices operating on battery-powered microprocessors or integrated circuits with limited computational resources and energy budget.

To implement real-time DL inference for clinical neural implants, three paradigms have been proposed (figure 1, left). First, inference can be done remotely through cloud computing (Paradigm A) [4]. In this paradigm, the neural implant transmits the recorded data to a wearable device via local wireless communication, and the wearable device in turn uploads the data to a cloud-based workstation via internet or telecommunication. The cloud-based inference result is downloaded wirelessly back to the wearable device and then to the neural implant for producing the control output (e.g. neural stimulation). Second, inference can be done on the wearable device without the need for data transfer to a cloud-based platform (Paradigm B) [10]. Third, DL inference can be performed directly on the neural implant itself, requiring no data transfer to another device (Paradigm C) [11]. Each of these three inference paradigms has its strengths and weaknesses (figure 1, right). Although

the cloud computing in Paradigm A offers the best possible inference accuracy, the latency and the robustness are the worst, which can adversely effect the timing-critical closed-loop therapeutic intervention. There are also concerns about cybersecurity and data privacy due to the required data transfer via internet or telecommunication [12]. Paradigm B eliminates the dependence on remote data transfer and the associated concerns, but it still relies on the robustness and security of the local wireless communication between the wearable devices and the neural implants. Moreover, the inference model complexity will be limited by the computational and power resources of the wearable devices. Paradigm C, using edge computing [13], avoids the disadvantages and concerns of wireless communication and can therefore potentially achieve the best robustness, the lowest latency, and the minimum security risks. Even if occasional wireless data transfer is performed in Paradigm C for offline analysis and performance assessments, the therapy itself would be robust to data transfer disruption and the opportunity for malicious attack would be the smallest of the three paradigms due to infrequency of transmission. However, a major question regarding Paradigm C is whether an edge DL model can achieve desired inference performance. Although machine learning-enabled processors for neural implants have been reported [14, 15], DL models have only rarely been implemented [11, 16, 17]. Therefore, the objective of the present work was to evaluate the inference performance limitations and resource constraints of edge DL designs through a case study.

For a DL model to be successfully deployed on a neural implant for real-time inference, the model design needs to meet three requirements. First, the model's inference performance (e.g. accuracy, sensitivity, specificity, etc) should meet the requirements of the target application, including potential degradation from model compression and quantization. Second, the model's inference time should meet the latency requirement, especially in closed-loop applications. This not only requires a sufficient computational speed, but also limits the maximum length of the buffered input data segment, resulting in a constraint on the model's architecture. Third, the power and resource costs of the DL model should meet the budget of the neural implant. For instance, the available memory resources, both non-volatile storage memory and random access memory (RAM), limit the total number of trainable model parameters.

Potential hardware platforms for neural implants with edge DL include dedicated AI accelerators and general-purpose low-power processors. Dedicated AI accelerators provide higher energy efficiency than general-purpose processors. However, commercially



available AI accelerators, such as IBM's TrueNorth [18], Google's Edge TPU [19], and Intel's Loihi [20], are too power hungry for neural implants. Also, large volume mobile devices-oriented AI intellectual property cores are not easily accessible for research-grade prototyping [21]. Ultra low-power AI accelerator chips have mainly been developed in research labs [22–24]. On the other hand, general-purpose low-power processors, such as ARM Cortex®, Texas Instruments MSP430®, and open-source reduced instruction set computer (RISC)-based microcontrollers (MCUs) provide a solution for low-cost, rapid prototyping for medical research and pre-clinical trials. If a further reduction in power or device footprint is desired, the MCU-based design can be migrated to an application-specific integrated circuit (ASIC) by integrating the MCU core together with analog neural recording and stimulation circuits, with an optional wireless communication module [14, 25]. Thus, in this work, we focused on edge DL design using a general-purpose MCU.

To investigate DL design and optimization methods fulfilling practical requirements of neural implants, we conducted a case study of epileptic seizure detection. Real-time seizure detection and intervention through closed-loop neural stimulation has proven to be an effective treatment for medically-refractory epilepsy [26]. Due to the importance of accurate, low-latency detection by a battery-powered neural implant, this application could benefit from edge computing. We adopted three commonly used DL architectures and customized the design for seizure detection on a MCU. We validated the methods using a publicly available annotated epilepsy database. The models were trained offline using Tensorflow. Compression and quantization methods were investigated for reducing the computational

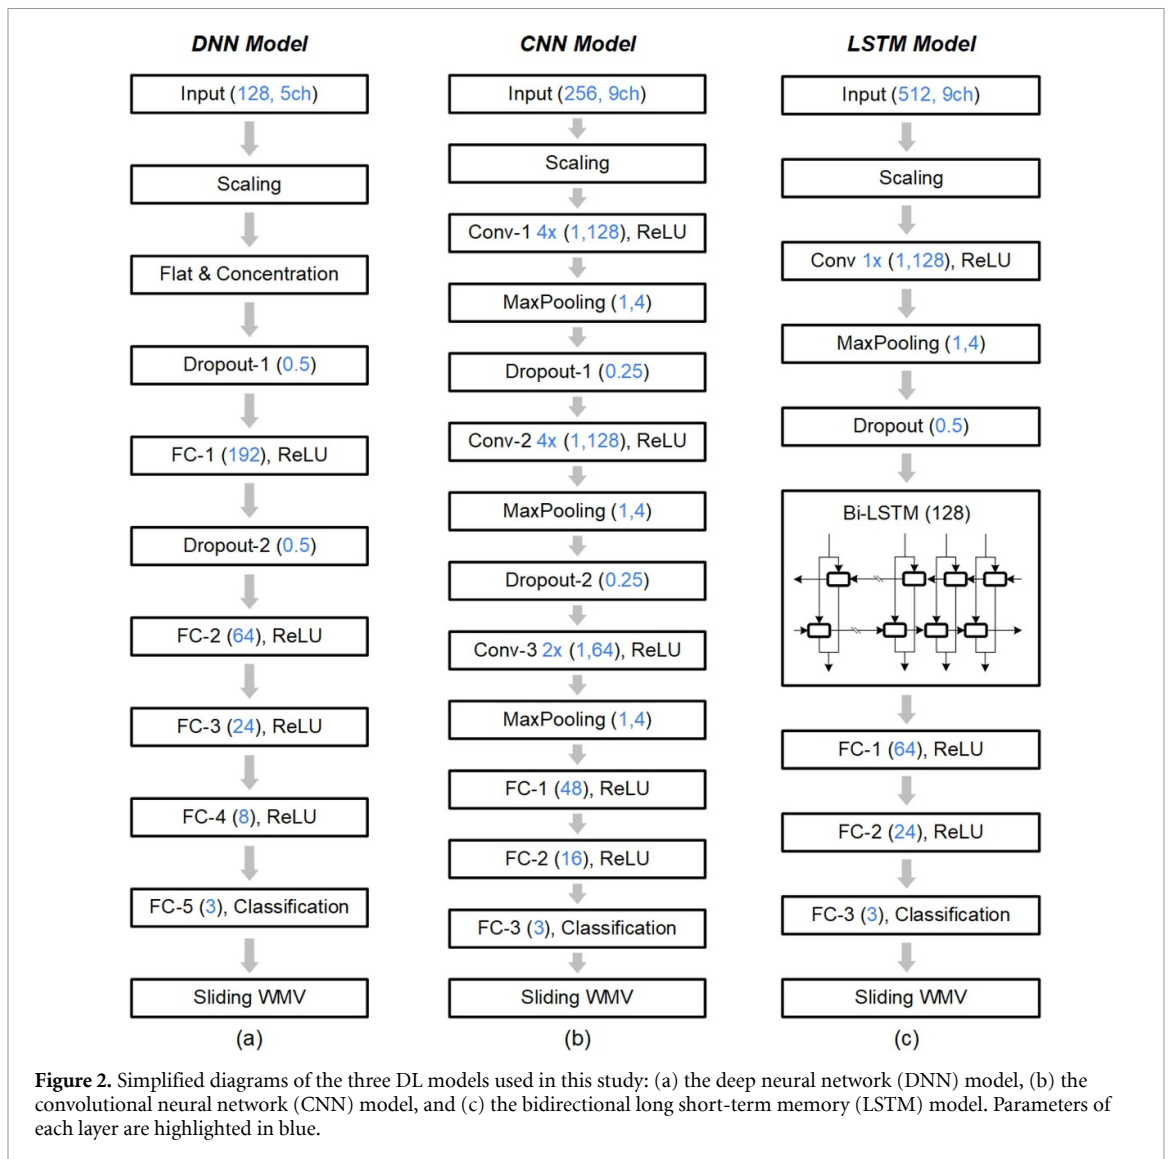
cost, memory, and power consumption for real-time inference. The inference results were compared to prior publications with no time and computational resource limitations. The strengths and weaknesses of each model were analyzed. Finally, we discuss the remaining challenges and future efforts toward implementing this novel paradigm in clinical neural implants.

## 2. Methods

### 2.1. Data preparation

This study used the Boston Children's Hospital (CHB)-MIT scalp electroencephalography (EEG) database [27], which is publicly available at PhysioNet.org [28]. Although scalp EEG is inconsistent with signals utilized by real neural implants, the CHB-MIT database is one of the most popular epilepsy data sets used for benchmarking, allowing a fair performance comparison between our edge DL approach and prior methods. The database contains recordings collected from 23 pediatric patients with intractable epilepsy. The placement of the surface electrodes followed the international 10–20 system [29]. The original recordings were sampled at 256 Hz and 16-bit resolution, with a 60 Hz notch filter applied for removing the mains interference. The seizure start and stop time points for each patient were manually annotated by clinical experts. Channels that were not constantly available throughout the entire duration in each case were excluded in developing the algorithms, as suggested in [9]. No signal processing was used before training and testing the models.

Data segments during ictal, preictal, and interictal phases were selected for training the DL models. The ictal phases were well defined by the expert



annotated seizure start and stop time points. There is no consistent definition of preictal phases in the literature, and the occurrence of preictal characteristics may differ across patients. In this work, we used the 3 min of recordings that are 30 s prior to each seizure's start time as the preictal phases. It should be noted that preictal characteristics could happen much earlier than this period [30]. Interictal segments were selected at least 2 h before and after any ictal phase to avoid potential signal contamination [31]. An even longer interval between ictal and interictal segments up to 4 h may further improve the classification performance [9, 32]. Unfortunately, the recordings of those phases are not always available in the chosen database. Postictal recordings were not used in this study.

The recordings from the interictal and preictal phases were much longer than the ictal phase. If all data segments were used in training, the severe imbalanced class distribution would cause a bias in the

model's prediction [33]. We adopted two methods to address this issue. First, we generated more ictal segments by sliding the window with an overlap of 50% [31]. Second, we applied a class weight to the loss function [3]. Interictal segments were randomly selected during day and night to avoid overfitting to irrelevant activities.

## 2.2. Deep learning models

We investigated three DL models for the seizure detection task: (a) a deep neural network (DNN) model, (b) a convolutional neural network (CNN) model, and (c) a bidirectional long short-term memory (LSTM) network model. The architectures of the three models are illustrated in figure 2. Although the model training was patient-specific, each model's architecture and parameter settings were kept the same for all patients in the database. The design details are presented in the following sections.

### 2.2.1. Deep neural network (DNN)

DNN is a type of feedforward artificial neural network that consists of multiple layers. Each layer contains many processing units, namely artificial neurons. The processing function of each neuron is given in the appendix. By connecting these neurons, DNNs emulate the way the brain processes information. Each neuron's output is processed by an activation function. In this work, the rectified linear unit (ReLU) function is used in all models [34]. Conventional artificial neural networks often use nonlinear activation functions, such as hyperbolic tangent and sigmoid. However, these functions make gradient-based training challenging as the number of network layers increases, an issue known as the gradient vanishing problem [35]. The introduction of the ReLU function overcomes this limitation by preserving the linear properties of positive values for gradient-based optimization, while still providing nonlinearity by setting all negative values to zero [34]. Moreover, the ReLU function is computational friendly, which is important for the purposes of this work.

Our DNN model consists of a scaling layer, a flat and concatenation layer, followed by five layers of fully-connected (FC) neurons. The scaling layer removes the DC offset from each channel and linearly scales the input data based on the dynamic range. The FC layers form a pyramid shape with  $N$  neurons in the first layer, where  $N$  is the length of the input segment of one channel. A 50% dropout layer is inserted before each of the first two FC layers. The class (ictal, preictal, or interictal) that has the highest activation function value is the final classification result.

### 2.2.2. Convolutional neural network (CNN)

CNNs have achieved extensive success in visual image recognition and natural language processing [36, 37]. Using CNNs in seizure detection tasks has also been reported [31, 38, 39]. Convolutional (Conv) layers can be trained to automatically extract underlying features that best represent the data without human intervention. Given their superior ability in extracting features from images, past studies have converted time-domain EEG segments into spectrograms to be used as the CNN inputs [31]. However, the computational costs of the Fourier or wavelet transform of multiple EEG channels are prohibitive for real-time applications using low-power MCUs. In this work, we used Conv layers to extract features directly from time-domain EEG signals.

Instead of using 2D standard Conv filters, we used 1D Conv filters to process each EEG channel. This allowed us to reduce the model size and computational cost, while still preserving the most important features [3]. Our CNN model consists of three Conv layers, all using 'same' padding and stride = 1. Each of the first two Conv layers consists of four

kernels with a length of  $fs/2$ , where  $fs$  is the sampling rate of the input EEG signal. The third Conv layer consists of two kernels with a length of  $fs/4$ . A max-pooling layer with a length of 4 was added to each of the Conv layers. The max pooling was used to prevent overfitting while reducing the computational costs. A 25% dropout layer was inserted after each of the first two Conv layers to regularize the model. The outputs of the Conv layers were concatenated and processed by three FC layers.

### 2.2.3. Long short-term memory network (LSTM)

LSTM is a type of recurrent neural network (RNN). In contrast to the feedforward neural networks described above, a RNN has recurrent connections that are suitable for capturing sequential information in the data. In particular, gating functions are used in each cell of the LSTM layers to control precisely what information is to be kept in the network and what is to be removed. Thus, LSTM has an inherent advantage in extracting certain temporal characteristics in time-domain signals [40], which is crucial in tasks such as seizure detection.

In this work, we adopted a bidirectional LSTM network. Bidirectional LSTM networks process sequential information from two opposite directions simultaneously, which has proven useful in applications such as speech processing [41]. For the gate activation function, we used hard-sigmoid instead of sigmoid to avoid the exponential operation, and we used softsign as the state activation function. The processing and activation functions of the LSTM model are given in the appendix.

We adopted a topology that combines an input Conv layer with the bidirectional LSTM layers. One kernel with a length of  $fs/2$  was used in the Conv layer, followed by a max-pooling layer with a length of 4. A 50% dropout layer was inserted between the input Conv layer and the bidirectional LSTM layers. Each of the bidirectional LSTM layers consists of 128 cells. Finally, FC layers convert the outputs for classification.

### 2.2.4. Sliding window based weighted majority voting

Short data segment based classification often suffers from a trade-off between sensitivity and FPR [4]. Since not every single data segment in the preictal phase exhibits signal characteristics that are related to an oncoming seizure, there is usually a limitation of the achievable classification accuracy before model overfitting. To achieve a high sensitivity while minimizing the FPR, we propose a novel sliding window based weighted majority voting (WMV) algorithm. The algorithm was implemented for real-time operation, as described in the pseudo-code of algorithm 1.



**Algorithm 1:** Sliding window based WMV algorithm.

---

**Inputs:**  $Pred\_Seg [1 : M] \in \{Ictal, Preictal, Interictal\}$ , ;  
 $M$ : total segments in the win,  
Parameters :  $\alpha_{I,P}$ ,  $\beta_{I,P}$ , and  $\theta_{I,P}$

**Output:**  $Pred\_Event \in \{Ictal, Preictal, Interictal\}$

- 1 **Initialization:**
- 2  $Score[Ictal] \leftarrow 0, Score[Preictal] \leftarrow 0,$   
 $Score[Interictal] \leftarrow 0$
- 3  $Acc[Ictal] \leftarrow 0, Acc[Preictal] \leftarrow 0$
- 4  $Pred\_Event \leftarrow Interictal$ ;
- 5 **for**  $i \leftarrow 1$  **to**  $M$  **do**
- 6     **switch** based on  $Pred\_Seg [i]$  **do**
- 7         **case**  $Ictal$  **do**
- 8              $Score[Ictal] \leftarrow + = \alpha_I + \beta_I \cdot Acc[Ictal]$
- 9              $Acc[Ictal] \leftarrow + = 1, 10$
- $Acc[Preictal] \leftarrow 0$
- 10         **end**
- 11         **case**  $Preictal$  **do**
- 12              $Score[Preictal] \leftarrow + = \alpha_P + \beta_P$
- $Acc[Preictal]$
- 13              $Acc[Preictal] \leftarrow + = 1, Acc[Ictal] \leftarrow 0$
- 14         **end**
- 15         **case**  $Interictal$  **do**
- 16              $Acc[Ictal] \leftarrow 0, Acc[Preictal] \leftarrow 0$
- 17         **end**
- 18     **end**
- 19     **if**  $Score[Ictal] > \theta_I$
- 20          $Pred\_Event \leftarrow Ictal$
- 21         **Break**;
- 22     **else if**  $Score[Preictal] > \theta_P$
- 23          $Pred\_Event \leftarrow Preictal$
- 24         **Break**;
- 25     **else**
- 26         **Continue**;
- 27     **end**
- 28 **end**

---

The algorithm uses the DL model's segment-based classification results as inputs. The evaluation is based on a sliding window of  $M$  segments. A score for the present window to be in an ictal phase is calculated based on two weighted terms: (a) how many individual segments generate a classification result of ictal phase, weighted by a coefficient of  $\alpha_I$ , and (b) how many times the ictal classification repeats in a row, weighted by a coefficient of  $\beta_I$ . The score for the present window to be in a preictal phase is estimated in an equivalent way, based on the two terms corresponding to preictal phases, weighted by coefficients  $\alpha_P$  and  $\beta_P$ . Finally, two pre-defined thresholds  $\theta_I$  and  $\theta_P$  are used to determine the final classification result from the sliding window. Once a score crosses the threshold, the algorithm will mark the event and break from evaluating the current sliding window. This is mainly to avoid delays during real-time seizure detection. A new evaluating sliding window starts immediately after the previous window terminates. Compared with traditional majority voting or moving average based algorithms, the proposed algorithm favors a prediction if the same classification results appear in a continuous manner. This significantly reduces the FPR in the classification.

Moreover, the detection latency is not limited by the sliding window length as the algorithm terminates once the thresholds are crossed.

### 2.3. Model compression and quantization

For the models to be successfully deployed in a MCU, we applied compression and quantization techniques to reduce the computational cost. There are existing channel pruning techniques for reducing model dimension without retraining with data [42]. However, since our model architectures are compact, the most effective method of compression is through iterations of retraining.

To select the most informative channels as the model inputs, we ranked all available recording channels based on the line length feature of these channels during the ictal phase. Line length is a measure of both high amplitude and high-frequency content of a time-domain signal, and is proven to be among the most effective features of seizures [43, 44]. It is computed as:

$$f_L(x_i) = \frac{1}{N} \sum_{t=1}^{N-1} |x_i(t-1) - x_i(t)| \quad (1)$$

where  $x_i(t)$  is the EEG signal of channel  $i$  at time  $t$ ,  $N$  is the sample count of the ictal segment. The line length feature was only computed offline for channel ranking. It was not used in real-time seizure detection. Based on the channel ranking, the top  $K$  channels were used as the inputs for the models. The classification performances using different  $K$  were compared for determining the optimal channel set. Similarly, we compared the performance using different data segment length  $N$ .

Computing a DL model using high arithmetic precision does not necessarily improve performance [45]. Quantization of high precision coefficients not only saves the computational cost, but also reduces the memory cost for storing the coefficients. In fact, loading the coefficients from memory to the arithmetic unit can dominate the energy consumption in a microprocessor [46]. In this work, we quantized all coefficients to 8-bit fixed-point numbers. To evaluate the effects of model quantization, the performance of 8-bit quantized models was compared to that of 16-bit models (referred to hereafter as unquantized models). Together with the computationally efficient non-linear activation functions (appendix equations (7), (14) and (15)) used in the model, the computational cost and latency were reduced.

### 2.4. Training and testing methods

The DL models were trained by supervised learning using the annotated data segments. Adam optimizer was used in training all of the models [47]. The experiments for model compression and quantization were performed in a 10-fold cross validation (CV). We

used stratified CV where each fold has equal number of instances for each class. The number of segments depends on the available ictal onset time in each subject's recording, and the data segment length. The total data segments used during the 10-fold CV were less than 1% of the whole recording per subject.

After finalizing the model architectures, we evaluated the performances using a variant of the leave-one-out cross validation (LOOCV). For a subject's recording that contains  $K$  seizure events, the recording is divided into  $K$  sections with one seizure event in each of them. One section is used for validation and the rest  $K - 1$  sections are used for training. The process is repeated  $K$  times so that all data is used exhaustively for validation. LOOCV suffers from large variation when  $K$  is small [48]. In this work, we evaluate cases in the CHB-MIT that have at least five seizures in the recording, which lead to 15 cases including chb01, 03, 05, 06, 08, 10, 12–16, 18, 20, 23 and 24. Unlike the segment based 10-fold CV, the LOOCV is performed in a continuous manner in real-time using the sliding WMV algorithm.

Standard classification metrics were used, including accuracy, sensitivity, specificity, and FPR:

$$\text{Sensitivity} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (2)$$

$$\text{Specificity} = \frac{\text{TN}}{\text{TN} + \text{FP}} \quad (3)$$

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \quad (4)$$

$$\text{FPR} (h^{-1}) = \frac{\text{FP}}{\text{Total recording length (in hours)}} \quad (5)$$

where TP, TN, FP, and FN are true positive, true negative, false positive, and false negative detection, respectively. During the segment-based evaluation, results were directly compared with the annotated labels. For event-based seizure detection, we made the following definitions:

- (a) If a detection event happens within 5 s of the seizure start time, it is considered as a TP. There can be at most one TP per genuine seizure.
- (b) If no detection event happens within 5 s of the seizure start time, it is considered as a FN. There can be at most one FN per genuine seizure.
- (c) If a detection event happens earlier than 5 s of the seizure start time, or 5 s after the seizure end time, it is considered to be a FP. The FP count is not limited by the number of seizures.

Similarly, we made the definitions for event-based seizure prediction as follows:

- (a) If a warning alert happens within 40 min before the seizure start time, it is considered as a TP. There can be at most one TP per genuine seizure.
- (b) If no warning alert happens within 40 min before the seizure start time, it is considered as a FN. There can be at most one FN per genuine seizure.
- (c) If a warning alert happens earlier than 40 min before or after the seizure start time, it is considered to be a FP. The FP count is not limited by the number of seizures.

MATLAB® was used for data handling. The DL models were implemented in Python with Tensorflow, which is an open-source machine learning library developed by the Google Brain team [49]. The DL training was performed on Google Cloud clusters using tensor processing units.

## 2.5. Hardware implementation

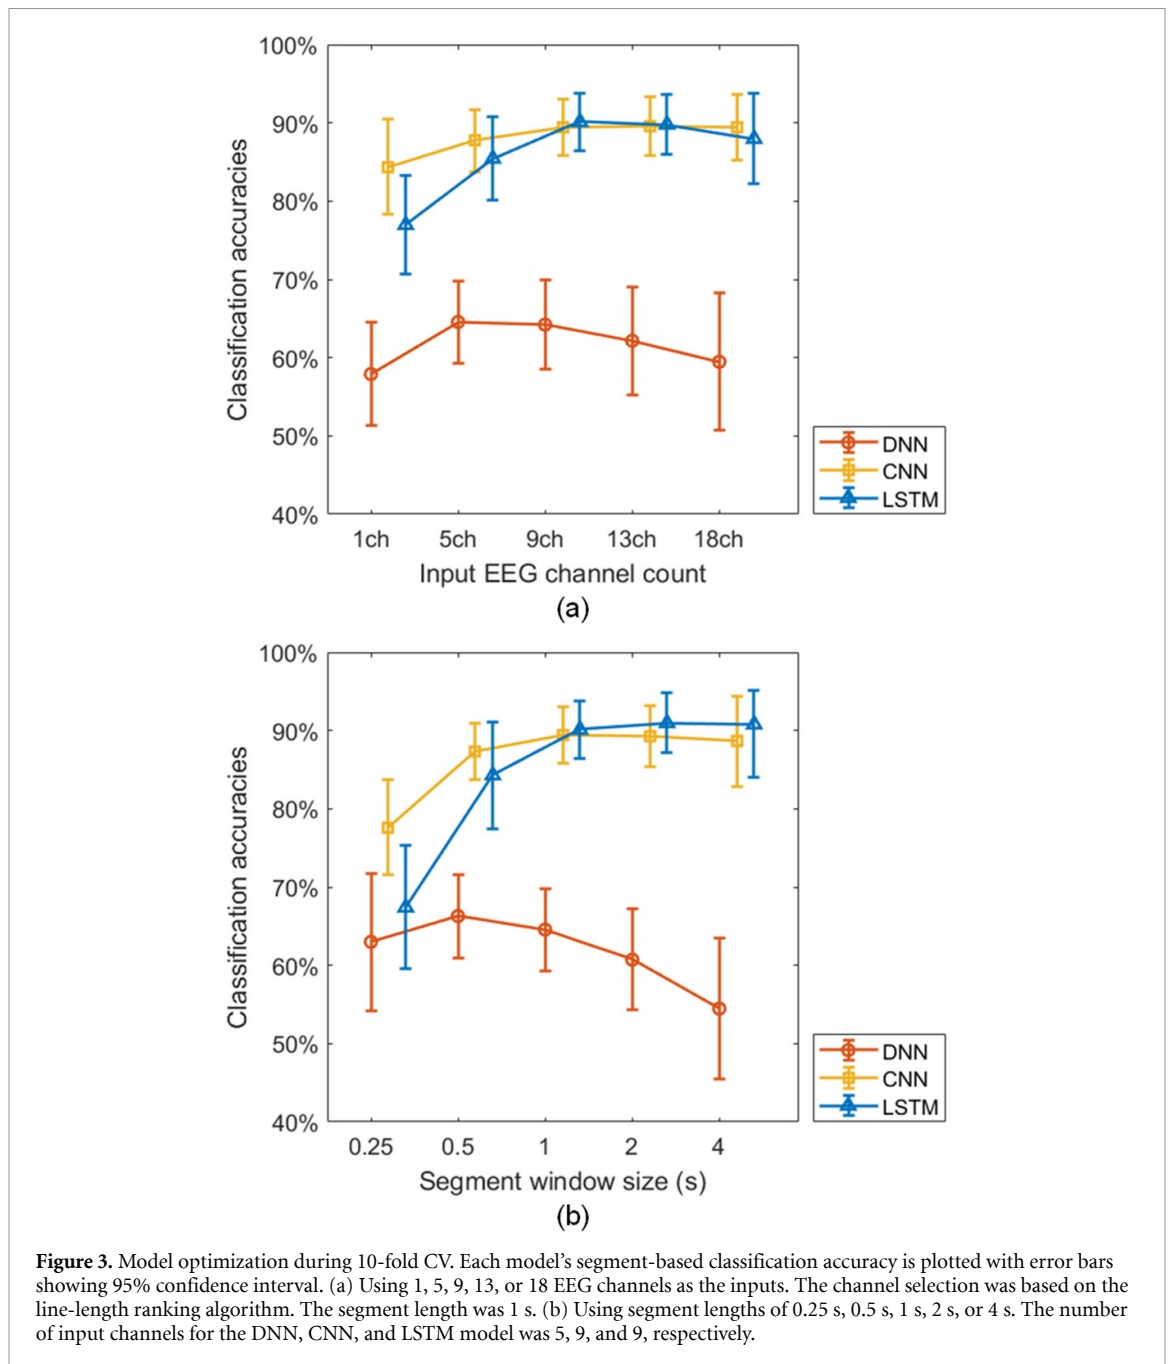
After training, the compressed and quantized models were deployed on a low-power MCU. We used a 32-bit ARM® Cortex-M4 based MCU nRF52840 from Nordic Semiconductor [50]. The nRF52840 features a flash memory size of 1 MB and a RAM size of 256 kB. The MCU runs at a clock rate of 64MHz. The Cortex-M4 core supports multiple types of hardware multiplication in one clock cycle, including 16-bit signed multiplication with 32-bit results [51]. The integrated floating-point unit was not used in this work.

The trained DL models were implemented in C/C++ for programming the MCU. The code was developed using Keil® MDK [52]. Open-source and commercial tools can be used to assist the code conversion from trained DL models to embedded systems [53, 54]. The neural signal acquisition process was not implemented in this work. Instead, the integrated USB 2.0 full speed module was used for transferring the recorded EEG data from the computer to the MCU. The input data segments were buffered in the RAM. The direct memory access module was used for transferring the data so that the CPU was not interrupted [55]. The inference results were returned to the computer via the same USB interface.

## 3. Results

### 3.1. Model optimization

Three DL models were trained on all 24 cases (from 23 patients) in the CHB-MIT database. Models were not quantized during the training and optimization phase. The input channels were selected among all available EEG channels for each patient using the proposed line-length based ranking (equation (1)). Figure 3(a) shows the classification performance of each model using 1, 5, 9, 13, or 18 channels as the inputs. The segment length was chosen to be 1 s (256 samples) in this analysis. The experimental result suggests that the optimal channel count is different for



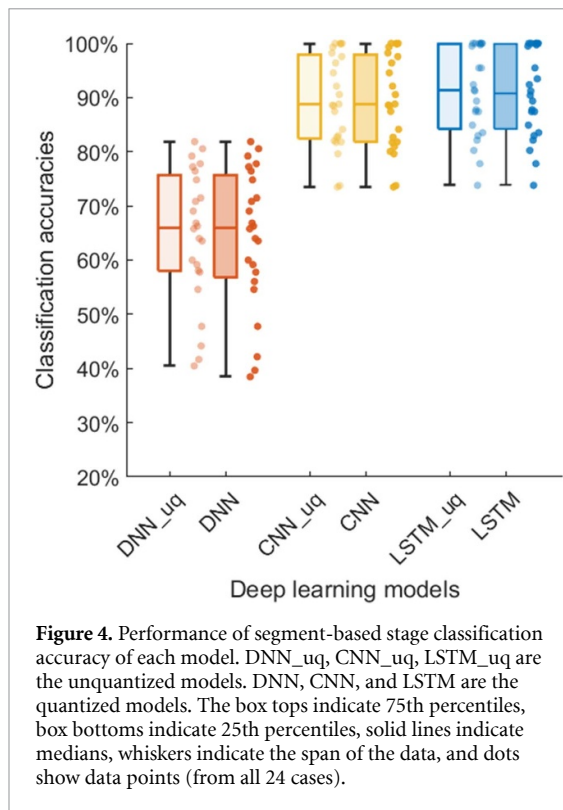
each model. Given the existing model size and experimental setup, the DNN model can take no more than five channels before it fails to extract critical information. More channels were generally helpful for the CNN model, but the performance improved marginally beyond nine channels. The LSTM model reached peak performance at 9 and 13 channels, while more channels only caused additional variance.

Segment length is another key parameter that impacts the overall performance as well as the computational cost. Figure 3 (b) shows the classification performance of each model using a segment length of 0.25 s, 0.5 s, 1 s, 2 s, or 4 s. The ictal segments were generated with 50% overlap in all cases. In the DNN model, using five channels as the input, the optimal performance was obtained with a segment length of

0.5 s. In the CNN model, using nine channels as the input, the optimal performance was obtained with a segment length of 1 s. In the LSTM model, also using nine channels as the input, the optimal performance was obtained with a segment length of 2 s.

We targeted the optimal performance of each model that the hardware resources permit. In certain model configurations, using more channels as inputs may reduce the segment size for achieving similar performance. Depending on the system's requirement and the available hardware resources, one may prefer to use more input channels or a larger segment size. In practice, more input channels require more wearable/implantable electrodes and corresponding recording hardware, such as low-noise neural amplifiers. On the other hand, a larger window size requires





more model coefficients thus more memory and RAM resources, and may cause a systematic latency. It should also be noted that increasing the segment length reduces the number of training and testing data segments, especially for patients with short ictal phases (e.g. 6-9 s in chb16). This could be a limiting factor in training DL models using a limited database. The hyperparameters used for training should be carefully tuned to minimize the impact.

Figure 4 shows the segment-based classification accuracy of each model. The DNN\_uq, CNN\_uq, and LSTM\_uq indicated in the figure are the 16-bit unquantized models, while DNN, CNN, and LSTM are the 8-bit quantized versions. The average performance degradation due to quantization was less than 0.5%. The LSTM model achieved the best overall performance (90.94%), followed by the CNN model (89.21%). The performance of the DNN model was relatively poor (64.55%).

To compare the classification performance with non-DL algorithms, we constructed a linear discriminant analysis (LDA) classifier. Spectral amplitudes in selected frequency bands were used as the input features. The LDA classifier was chosen because it has been used in the Medtronic Activa PC+S device [4]. The selected frequency bands were 0–2.7 Hz, 2.7–5.4 Hz, 5.4–10.8 Hz, 10.8–21.7 Hz, 21.7–43.4 Hz, and 43.4–86.8 Hz [56]. Table 1 shows the segment-based classification accuracy, sensitivity, and specificity of the LDA classifier and the three DL models for 10-fold CV. The DL models achieved better performance than the LDA classifier. Among the three DL models, the CNN and LSTM models showed superior overall

performance, while the DNN model had a limited ability in discriminating the preictal and ictal phases.

### 3.2. Model performance evaluation

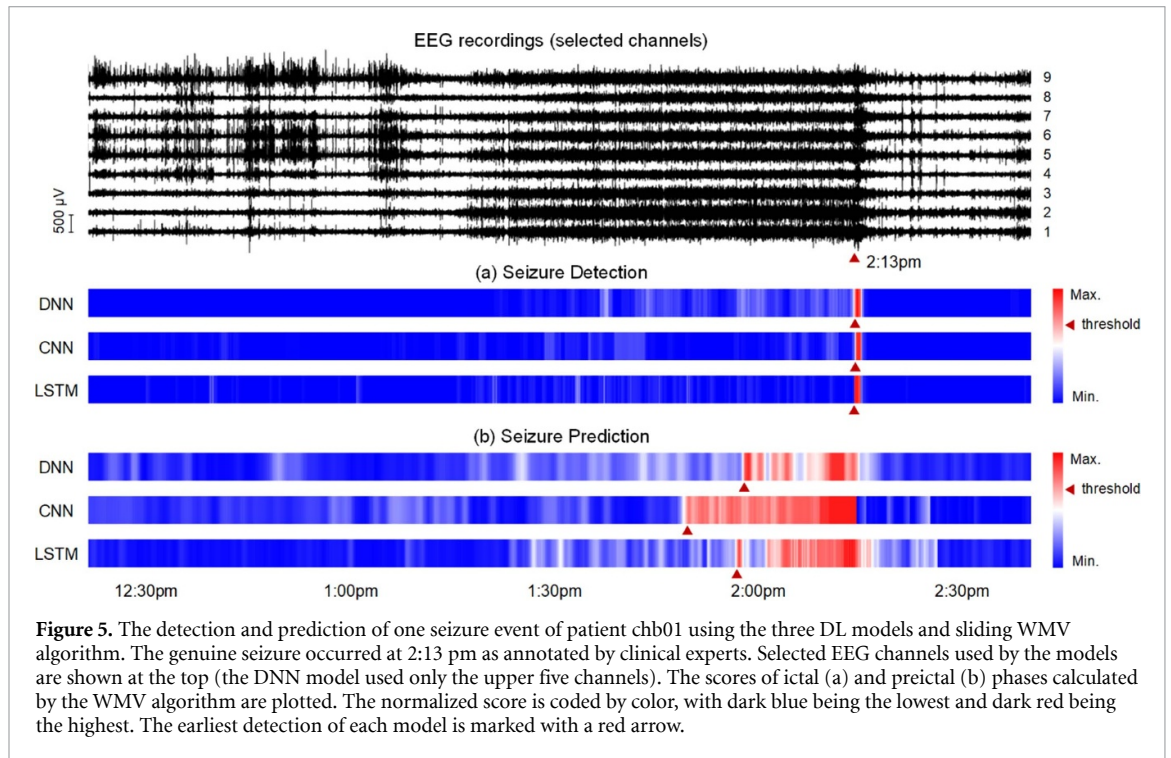
Real-time seizure event detection and prediction using the trained, quantized DL models were simulated by streaming the selected EEG time series to the MCU in a continuous manner. Figure 5 shows an illustrative example of the inferred scores of ictal ( $Score [Ictal]$  in algorithm 1) and preictal ( $Score [Preictal]$  in algorithm 1) phases calculated by each model at each instance in time. The selected input EEG channels for the CNN and LSTM models are shown at the top of figure 5. Only the upper five channels were used as inputs to the DNN model. All three models successfully detected the seizure event within a 5 s window around the marked start time (figure 5(a)). Furthermore, all three models successfully predicted the seizure within a 40 min horizon prior to the actual onset time (figure 5(b)). In this example, the CNN model showed the best robustness against false positives in predicting seizures.

The performance of event-based seizure detection and prediction was tested with LOOCV. The performance of the WMV algorithm was compared with a traditional moving average. The WMV algorithm improved the seizure detection FPR from  $0.745 \text{ h}^{-1}$  to  $0.169 \text{ h}^{-1}$  and the seizure prediction FPR from  $2.341 \text{ h}^{-1}$  to  $0.710 \text{ h}^{-1}$ . Figure 6 shows the box chart of each model's performance in seizure detection and prediction tasks in LOOCV. For seizure detection, the LSTM model achieved the highest average sensitivity (97.61%) and the lowest FPR ( $0.071 \text{ h}^{-1}$ ), which corresponds to one false alarm in every 14.1 h. For seizure prediction, the LSTM and CNN models achieved a comparable sensitivity above 90%. The CNN model had a better average FPR ( $0.204 \text{ h}^{-1}$ ), which corresponds to one false alarm in every 4.9 h. The performance of the DNN model was relatively poor in seizure prediction tasks. The average and median sensitivities and FPRs of each model are summarized in table 2.

Finally, each model's memory size, inference execution time, and power consumption are shown in figure 7. The results before and after quantization were plotted for comparison. The memory size reflects the actual hardware implementation including the code overhead for data handling. The data transfer time via the USB port was excluded from the inference time since this would not be present in autonomous neural implants. The CPU core was put in sleep mode after executing the inference, and the power consumption was measured directly from the power supply. The reported power consumption is an average within the segment period. The DNN model required the largest memory size, while the CNN model required the least. This is mainly because of the limited convolutional kernel size used in the CNN model. The LSTM model had the longest inference time, but it was still within its segment

**Table 1.** Segment-based classification in 10-fold CV using quantized models.

		LDA	DNN	CNN	LSTM
Overall accuracy		62.07%	64.55%	89.21%	90.94%
Sensitivity	Ictal	83.67%	82.52%	96.59%	97.30%
	Preictal	49.24%	51.83%	88.22%	91.46%
	Interictal	55.53%	60.66%	83.67%	85.33%
	Avg.	62.81%	65.00%	89.50%	91.53%
Specificity	Ictal	83.05%	88.89%	97.79%	98.68%
	Preictal	73.55%	75.29%	90.47%	91.58%
	Interictal	74.94%	83.75%	96.32%	97.00%
	Avg.	77.18%	82.64%	94.86%	95.75%



**Figure 5.** The detection and prediction of one seizure event of patient chb01 using the three DL models and sliding WMV algorithm. The genuine seizure occurred at 2:13 pm as annotated by clinical experts. Selected EEG channels used by the models are shown at the top (the DNN model used only the upper five channels). The scores of ictal (a) and preictal (b) phases calculated by the WMV algorithm are plotted. The normalized score is coded by color, with dark blue being the lowest and dark red being the highest. The earliest detection of each model is marked with a red arrow.

period. From the perspective of power consumption, the DNN and CNN models were comparable, while the LSTM model consumed the most.

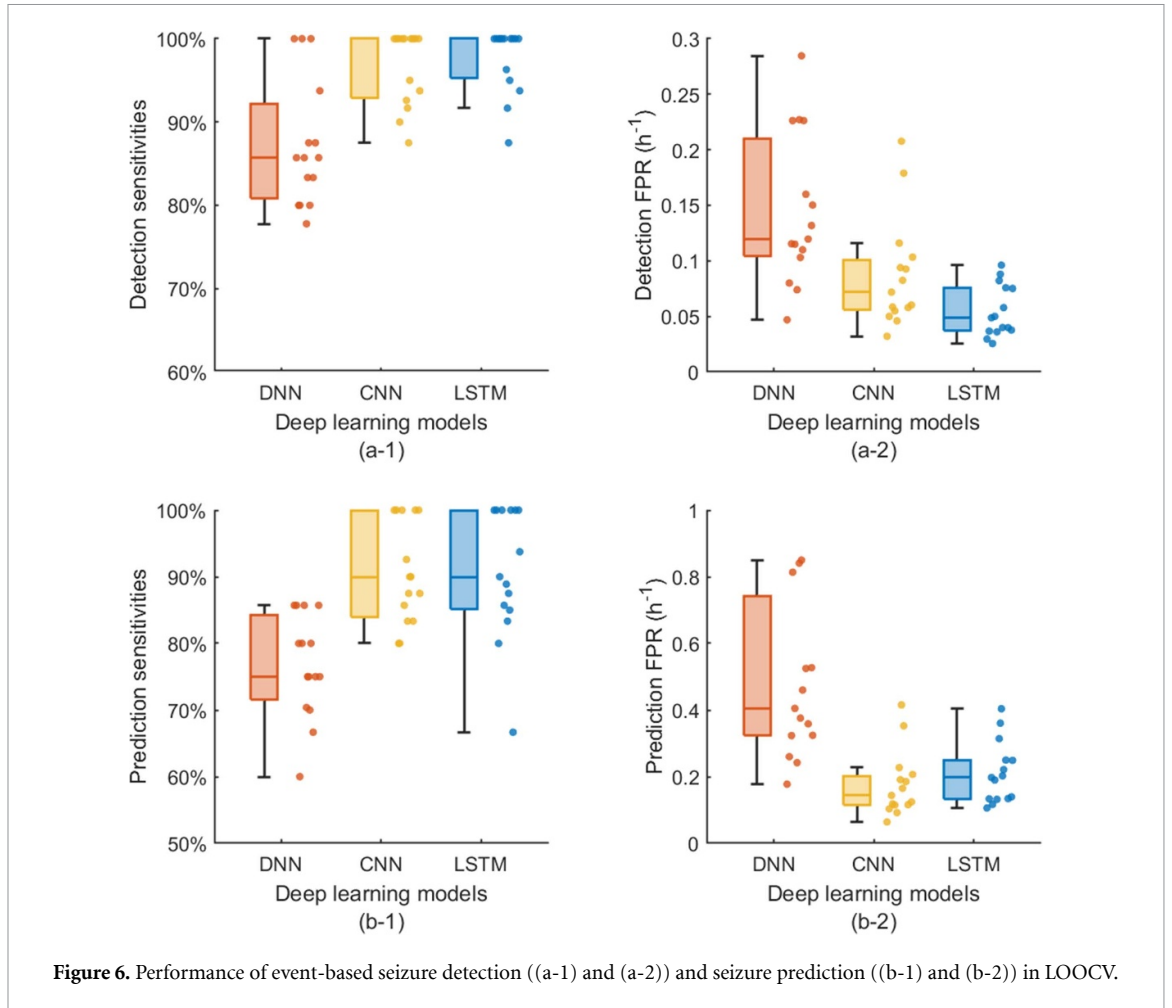
#### 4. Discussion

Each of the three DL models developed in this work has its own strengths and weaknesses. The DNN model achieved the shortest inference time with minimum power consumption. The CNN model achieved a balanced performance with moderate power consumption and the smallest memory cost. The LSTM model achieved the best overall performance (e.g. highest seizure detection sensitivity and lowest FPR) at the expense of relatively long inference time and high power consumption. The optimal choice of model mainly depends on the specific application as well as the available hardware resources of the neural implant, including the batteries.

Model quantization improved inference time, memory size, and power consumption without sacrificing inference performance. On the other hand,

model compression and channel pruning may require iterations of retraining. Layer size should be scaled with the sampling rate for capturing temporal features. Cascading DL models with a second stage sliding window-based algorithm, such as the sliding window based WMV algorithm implemented in this work, may compensate for the limitation of small input data buffer size of the edge DL models. Optimizing the design jointly for available hardware resources and real-time operational requirements is the key to achieving a satisfactory overall performance.

For event-based seizure detection, the CNN and LSTM models achieved an average high sensitivity of 96.7% and 97.61%, and low FPR at 0.102 and 0.071, respectively. The results suggest that they are potential candidates for real-time, closed-loop therapeutic systems. For seizure prediction, the performance of the DNN model was limited. We compare the seizure prediction performance of this work with prior studies (table 3). Importantly, most of the listed studies assumed no time and computational resource



**Table 2.** Performance of event-based seizure detection and prediction in LOOCV.

	Detection				Prediction			
	Sensitivity		FPR ( $h^{-1}$ )		Sensitivity		FPR ( $h^{-1}$ )	
	Avg.	Median	Avg.	Median	Avg.	Median	Avg.	Median
DNN	87.36%	85.71%	0.169	0.140	76.66%	75.00%	0.710	0.474
CNN	96.70%	100%	0.102	0.084	90.66%	90.00%	0.204	0.168
LSTM	97.61%	100%	0.071	0.063	90.72%	90.00%	0.241	0.227

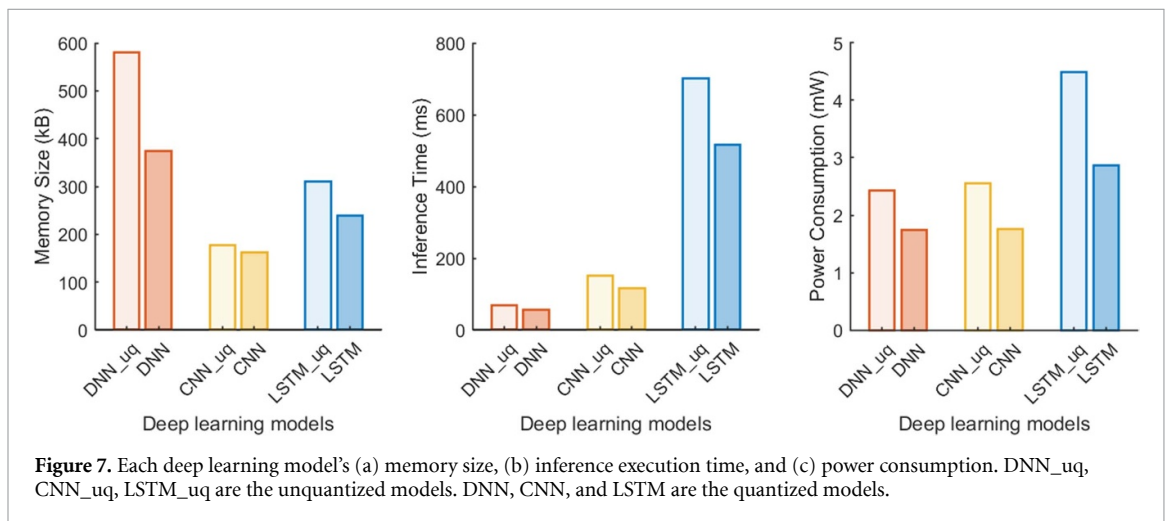


Table 3. Comparison with prior seizure prediction studies.

Year	Publication	Database	Features	Algorithm	Sensitivity	FPR ( $\text{h}^{-1}$ )	Prediction duration
2013	Li <i>et al</i> [65]	Freiburg (21 cases)	Spike rate	Threshold	72.7%	0.11	50 min
2016	Zhang <i>et al</i> [66]	CHB-MIT (17 cases)	PSD ratio	SVM	98.68%	0.047	50 min
2017	Chu <i>et al</i> [67]	CHB-MIT (13 cases)	Fourier Transform, PSD	Threshold	83.33%	0.392	86 min
2017	Alotaiby <i>et al</i> [68]	CHB-MIT (24 cases)	Spatial pattern statistics	LDA	81%	0.47	60 min
2017	Arabi <i>et al</i> [69]	Freiburg (10 cases)	Uni-/bivariate features	Rule-based	86.7%	0.126	30 min
2018	Khan <i>et al</i> [38]	CHB-MIT (15 cases)	Wavelet Transform	CNN	87.80%	0.142	10 min
2018	Truong <i>et al</i> [31]	CHB-MIT (13 cases)	Short-time Fourier Transform	CNN	81.20%	0.16	30 min
2018	Tsiouris <i>et al</i> [9]	CHB-MIT (24 cases)	Wavelet Transform, PSD, statistics, etc	LSTM	99.60%	0.006	30 min
2018	Shahbazi <i>et al</i> [70]	CHB-MIT (14 cases)	Short-time Fourier Transform	LSTM	98.2 %	0.13	45 min
2019	Affes <i>et al</i> [71]	CHB-MIT (24 cases)	Fourier Transform	C-GRNN	89.0%	1.6	35 min
2020	Abiyev <i>et al</i> [39]	CHB-MIT (7 cases)	Raw EEG	CNN	97.67%	1.97	N/A
	This work	CHB-MIT (15 cases)	Raw EEG	DNN	76.66%	0.710	40 min
				CNN	90.66%	0.204	40 min
				LSTM	90.72%	0.241	40 min

limitations. The highest performances were achieved at the cost of sophisticated, hand-crafted, patient-specific feature engineering [9]. However, such complicated processing typically prevents its application in embedded devices for real-time, closed-loop inference. The relatively good hardware-based seizure prediction performance of our CNN and LSTM models compared to prior software implementations provides support for their use in future therapies.

One caveat to seizure prediction is that more advanced notice is not always better, unless a seizure prediction horizon (SPH), which is a seizure-free warning period between the alarm and the actual seizure onset, can be guaranteed by the algorithm [57]. Otherwise, an early warning may increase the anxiety of the patient given that a seizure may or may not happen any time within a long period after the warning. Including an accurate SPH significantly increases the difficulty in prediction [57]. It is not applicable for the light-weight DL models developed in this work. It should be noted, however, that the proposed edge DL inference paradigm does not preclude the functionality of uploading data to an external system including cloud-based computing resources for further analysis. Periodical diagnosis and model updating may be necessary for clinical adoption. But since these operations are not continuous during everyday use, data encryption can be reinforced and the peaking power dissipation during these short periods is not a big concern.

Another caveat to the results is that the EEG signals in the selected data set were recorded noninvasively using scalp electrodes. Therapeutic neural implant devices typically acquire EEG signals intracranially with a much higher signal-to-noise ratio (SNR) [58]. The performance of DL algorithms typically improves with higher SNR signals, since more subtle neural features can be unveiled in these recordings [31]. We expect the performance of the models can be further improved, or the size and power consumption can be reduced if intracranial EEG recordings are used as the input.

Although this work uses a general-purpose, off-the-shelf MCU as the edge hardware platform, the design and optimization methods are applicable to ASIC development for clinical neural implants. The MCU core can be directly integrated into an ASIC design with the required memory or the models can be directly synthesized in the register-transfer level for minimizing the design overhead. Furthermore, in-memory or near memory computational techniques can be used to further reduce the power consumption of repeated multiply-accumulate (MAC) operations [22, 23]. An ASIC system-on-chip (SoC) that integrates analog neural interfaces (e.g. neural recorders and stimulators), digital DL inference module, power management and wireless communication modules can achieve the optimal power consumption and device footprint for chronic neural implants.

Finally, neural implants with different clinical purposes use various types of input neural signals, inference model complexities, and control strategies. Our work may serve as a reference for related studies on cognitive monitoring, sleep interventions, BMI, and other applications with closed-loop neuromodulation or neuroprosthetic control. The generalizability of our edge DL approach depends on the model complexity required to achieve high performance in these different applications. While DL model complexity has been defined in different ways, it is specifically the speed of inference, along with memory requirements, that is critical for resource-restrained edge devices [59]. To assess generalizability, we sought to compare our models' inference speed with that of a sample of prior DL studies in different neural application domains. Elapsed inference time (figure 7(b)) is rarely reported and is hardware dependent. A hardware-independent measure is the number of required computations (i.e. MAC operations) based on model architecture [60]. Restricting the comparison to CNN models, which are used in the majority of neural applications [5], the total number of operations in the Conv and FC layers provide a reasonable estimate of complexity (see appendix). Our CNN model required approximately 2.4 M MACs to process the  $256 \times 9$  input matrix. This complexity is higher than many prior CNN models used for BMI applications like P300 detection [61], steady-state evoked potential classification [62], and attentive state detection [63], which ranged from about 0.1 M to 0.5 M MACs. Although CNN models for other tasks like sleep scoring [64] can be significantly more complex, on the order of 10 M to 100 M MACs, we suggest that edge DL systems are a realistic option for many BMI applications. Since DL eliminates the domain-specific manual feature selection used in conventional algorithms, research progress and technological advancement in one specific application area should readily generalize to other applications.

## 5. Conclusion

In this work, we developed edge DL models and investigated their potential utility for future clinical applications involving neural implants. We adopted three commonly used DL architectures (DNN, CNN, and LSTM) and optimized the models for deployment in resource-restrained hardware. Using the CHB-MIT database, we show that edge DL inference can achieve comparable performance in epileptic seizure detection to many prior implementations that had no time and computational resource limitations. Our results suggest that edge DL inference is a promising option for closed-loop neuromodulation, with superior robustness and security compared to wireless communication-based solutions. While clinical studies are needed to confirm the efficacy of this



paradigm, this work demonstrates the feasibility and potential advantages. We envision that the next generation of clinical neural implants could leverage edge DL inference to greatly improve their therapeutic benefit.

### Data availability statement

No new data were created or analysed in this study.

### Appendix A. Processing and activation functions

In the FC layers, the processing function of each neuron  $j$  is given by:

$$y_j = f_R \left( \sum_{i=1}^n w_{j,i} \cdot x_i + b_j \right) \quad (6)$$

where  $x_i$  is from the previous layer,  $w_{j,i}$  is the weight factor,  $b_j$  is the bias term,  $n$  is the number of neurons in the previous layer, and  $f_R(x)$  is the activation function. The  $f_R(x)$  used in this work is the rectified linear unit (ReLU) function. ReLU is a piecewise linear function as given by:

$$f_R(x) = \max\{0, x\}. \quad (7)$$

The processing functions of each LSTM cell  $j$  are given by:

$$f_j = f_G(W_{h,f,j} \cdot h_{j-1} + W_{x,f,j} \cdot x_j + b_{f,j}) \quad (8)$$

$$i_j = f_G(W_{h,i,j} \cdot h_{j-1} + W_{x,i,j} \cdot x_j + b_{i,j}) \quad (9)$$

$$o_j = f_G(W_{h,o,j} \cdot h_{j-1} + W_{x,o,j} \cdot x_j + b_{o,j}) \quad (10)$$

$$\tilde{c}_j = f_S(W_{h,c,j} \cdot h_{j-1} + W_{x,c,j} \cdot x_j + b_{c,j}) \quad (11)$$

$$c_j = f_j \cdot c_{j-1} + i_j \cdot \tilde{c}_j \quad (12)$$

$$h_j = o_j \cdot f_S(c_j) \quad (13)$$

where  $x_j$  is the output from the previous layer,  $h_j$  is the hidden state,  $c_j$  is the memory state ( $\tilde{c}_j$  is the candidate),  $i_j$  is the input gate,  $o_j$  is the output gate, and  $f_j$  is the forget gate.  $W$ 's and  $b$ 's are the weight and bias terms for each neuron in the LSTM cell.

The gate activation function  $f_G(x)$  is given by:

$$f_G(x) = \max \left\{ 0, \min \left\{ 1, \frac{x}{5} + \frac{1}{2} \right\} \right\}. \quad (14)$$

The state activation function  $f_S(x)$  is given by:

$$f_S(x) = \frac{x}{1 + |x|}. \quad (15)$$

### A.1. Multiply-accumulate operation estimates

The number of MAC operations for a CNN can be estimated as follows. For FC layers:

$$\#MAC_{fc} = C_{in} \times C_{out} \quad (16)$$

where  $C_{in}$  is the number of input channels (neurons) and  $C_{out}$  is the number of output channels. For Conv layers:

$$\#MAC_{conv} = C_{in} \times m_{out} \times n_{out} \times C_{out} \times h \times v \quad (17)$$

where  $m_{out} \times n_{out}$  is the size of the output feature map,  $C_{out}$  is again the number of output channels (equivalent to the number of kernels), and  $h \times v$  is the size of each kernel. Max pooling and activation operations are ignored as they typically contribute very little to the total MAC count.

### ORCID iD

Xilin Liu  <https://orcid.org/0000-0002-9547-3905>

### References

- [1] Sun F T and Morrell M J 2014 Closed-loop neurostimulation: the clinical experience *Neurotherapeutics* **11** 553–63
- [2] Mak J N and Wolpaw J R 2009 Clinical applications of brain-computer interfaces: current state and future prospects *IEEE Rev. Biomed. Eng.* **2** 187–99
- [3] Lawhern V J, Solon A J, Waytowich N R, Gordon S M, Hung C P and Lance B J 2018 Eegnet: a compact convolutional neural network for EEG-based brain-computer interfaces *J. Neural Eng.* **15** 056013
- [4] Baldassano S, Zhao X, Brinkmann B, Kremen V, Bernabei J, Cook M, Denison T, Worrell G and Litt B 2019 Cloud computing for seizure detection in implanted neural devices *J. Neural Eng.* **16** 026016
- [5] Roy Y, Banville H, Albuquerque I, Gramfort A, Falk T H and Faubert J 2019 Deep learning-based electroencephalography analysis: a systematic review *J. Neural Eng.* **16** 051001
- [6] LeCun Y, Bengio Y and Hinton G 2015 Deep learning *Nature* **521** 436–44
- [7] Lotte F, Bougrain L, Cichocki A, Clerc M, Congedo M, Rakotomamonjy A and Yger F 2018 A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update *J. Neural Eng.* **15** 031005
- [8] Schirrmester R T, Springenberg J T, Fiederer L D J, Glasstetter M, Egensperger K, Tangermann M, Hutter F, Burgard W and Ball T 2017 Deep learning with convolutional neural networks for EEG decoding and visualization *Human Brain Mapp.* **38** 5391–420
- [9] Tsiouris K M, Pezoulas V C, Zervakis M, Konitsiotis S, Koutsouris D D and Fotiadis D I 2018 A long short-term memory deep learning network for the prediction of epileptic seizures using EEG signals *Comput. Biol. Med.* **99** 24–37
- [10] Mahmood M, Mzurikwao D, Kim Y S, Lee Y, Mishra S, Herbert R, Duarte A, Ang C S and Yeo W H 2019 Fully portable and wireless universal brain-machine interfaces enabled by flexible scalp electronics and deep learning algorithm *Nat. Mach. Intell.* **1** 412–22
- [11] Kiral-Kornek I *et al* 2018 Epileptic seizure prediction using big data and deep learning: toward a mobile system *EBioMedicine* **27** 103–11
- [12] Naufel S, Knaack G L, Miranda R, Best T K, Fitzpatrick K, Emondi A A, Van Gieson E and McClure-Begley T 2020

- Darpa investment in peripheral nerve interfaces for prosthetics, prescriptions and plasticity *J. Neurosci. Methods* **332** 108539
- [13] Hartmann M, Hashmi U S and Imran A 2019 Edge computing in smart health care systems: review, challenges and research directions *Trans. Emerg. Telecommun. Technol.* **e3710**
- [14] O'Leary G, Groppe D M, Valiante T A, Verma N and Genov R 2018 Nurip: neural interface processor for brain-state classification and programmable-waveform neurostimulation *IEEE J. Solid-State Circuits* **53** 3150–62
- [15] Zhu B, Shin U and Shoaran M 2020 Closed-loop neural interfaces with embedded machine learning *27th IEEE Int. Conf. on Electronics, Circuits and Systems (ICECS)* (Glasgow, UK) (<https://doi.org/10.1109/ICECS49266.2020.9294844>)
- [16] Heller S, Hügler M, Nematollahi I, Manzouri F, Dümpelmann M, Schulze-Bonhage A, Boedecker J and Woias P 2018 Hardware implementation of a performance and energy-optimized convolutional neural network for seizure detection *2018 40th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society (EMBC)* (IEEE) pp 2268–71
- [17] Hügler M, Heller S, Watter M, Blum M, Manzouri F, Dümpelmann M, Schulze-Bonhage A, Woias P and Boedecker J 2018 Early seizure detection with an energy-efficient convolutional neural network on an implantable microcontroller *2018 Int. Conf. on Neural Networks (IJCNN)* (IEEE) pp 1–7
- [18] Akopyan F *et al* 2015 Truenorth: design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **34** 1537–57
- [19] Cass S 2019 Taking AI to the edge: Google's TPU now comes in a maker-friendly package *IEEE Spectr.* **56** 16–7
- [20] Davies M *et al* 2018 Loihi: a neuromorphic manycore processor with on-chip learning *IEEE Micro* **38** 82–99
- [21] Wang X, Han Y, Wang C, Zhao Q, Chen X and Chen M 2019 In-edge AI: intelligentizing mobile edge computing, caching and communication by federated learning *IEEE Netw.* **33** 156–65
- [22] Si X *et al* 2019 A twin-8T SRAM computation-in-memory unit-macro for multibit CNN-based AI edge processors *IEEE J. Solid-State Circuits* **55** 189–202
- [23] Jia H, Valavi H, Tang Y, Zhang J and Verma N 2020 A programmable heterogeneous microprocessor based on bit-scalable in-memory computing *IEEE J. Solid-State Circuits* **2609–21**
- [24] Chen J and Ran X 2019 Deep learning with edge computing: a review *Proc. IEEE* **107** 1655–74
- [25] Liu X, Zhang M, Richardson A G, Lucas T H and Van der Spiegel J 2016 Design of a closed-loop, bidirectional brain machine interface system with energy efficient neural feature extraction and pid control *IEEE Trans. Biomed. Circuits Syst.* **11** 729–42
- [26] Heck C N *et al* 2014 Two-year seizure reduction in adults with medically intractable partial onset epilepsy treated with responsive neurostimulation: final results of the RNS system pivotal trial *Epilepsia* **55** 432–41
- [27] Shoeb A H 2009 Application of machine learning to epileptic seizure onset detection and treatment *PhD Thesis* Massachusetts Institute of Technology
- [28] Goldberger A L *et al* 2000 Physiobank, physiotoolkit and physionet: components of a new research resource for complex physiologic signals *Circulation* **101** e215–20
- [29] Homan R W, Herman J and Purdy P 1987 Cerebral location of international 10–20 system electrode placement *Electroencephalogr. Clin. Neurophysiol.* **66** 376–82
- [30] Litt B and Echauz J 2002 Prediction of epileptic seizures *Lancet* **27** 421–4
- [31] Truong N D, Nguyen A D, Kuhlmann L, Bonyadi M R, Yang J, Ippolito S and Kavehei O 2018 Convolutional neural networks for seizure prediction using intracranial and scalp electroencephalogram *Neural Netw.* **105** 104–11
- [32] Daoud H and Bayoumi M A 2019 Efficient epileptic seizure prediction based on deep learning *IEEE Trans. Biomed. Circuits Syst.* **13** 804–13
- [33] Japkowicz N and Stephen S 2002 The class imbalance problem: a systematic study *Intell. Data Anal.* **6** 429–49
- [34] Nair V and Hinton G E 2010 Rectified linear units improve restricted boltzmann machines *Proc. of the 27th Int. Conf. on Mach. Learn.* 807–14
- [35] Goodfellow I, Bengio Y, Courville A and Bengio Y 2016 *Deep Learning* vol 1 (Cambridge: MIT Press) (<https://doi.org/10.4258/hir.2016.22.4.351>)
- [36] Krizhevsky A, Sutskever I and Hinton G E 2012 Imagenet classification with deep convolutional neural networks *Advances in Neural Information Processing Systems* pp 1097–105
- [37] Kim Y 2014 Convolutional neural networks for sentence classification (arXiv:1408.5882)
- [38] Khan H, Marcuse L, Fields M, Swann K and Yener B 2018 Focal onset seizure prediction using convolutional networks *IEEE Trans. Biomed. Eng.* **65** 2109–18
- [39] Abiyev R, Arslan M, Idoko J B, Sekeroglu B and Ilhan A 2020 Identification of epileptic EEG signals using convolutional neural networks *Appl. Sci.* **10** 4089
- [40] Gers F A, Schmidhuber J and Cummins F 1999 Learning to forget: continual prediction with LSTM *Neural Comput.* **850–5**
- [41] Graves A, Jaitly N and Mohamed A R 2013 Hybrid speech recognition with deep bidirectional LSTM *2013 IEEE Workshop on Automatic Speech Recognition and Understanding* (IEEE) pp 273–8
- [42] He Y, Zhang X and Sun J 2017 Channel pruning for accelerating very deep neural networks *Proc. IEEE Int. Conf. on Computer Vision* pp 1389–97
- [43] Esteller R, Echauz J and Tcheng T 2004 Comparison of line length feature before and after brain electrical stimulation in epileptic patients *The 26th Annual Int. Conf. IEEE Engineering in Medicine and Biology Society* vol 2 (IEEE) pp 4710–13
- [44] Logesparan L, Casson A J and Rodriguez-Villegas E 2012 Optimal features for online seizure detection *Med. Biol. Eng. Comput.* **50** 659–69
- [45] Han S, Mao H and Dally W J 2015 Deep compression: compressing deep neural networks with pruning, trained quantization and Huffman coding (arXiv:1510.00149)
- [46] Le Gallo M, Sebastian A, Mathis R, Manica M, Gieffers H, Tuma T, Bekas C, Curioni A and Eleftheriou E 2018 Mixed-precision in-memory computing *Nat. Electron.* **1** 246–53
- [47] Kingma D P and Ba J 2014 Adam: a method for stochastic optimization (arXiv:1412.6980)
- [48] Kearns M and Ron D 1999 Algorithmic stability and sanity-check bounds for leave-one-out cross-validation *Neural Comput.* **11** 1427–53
- [49] Abadi M *et al* 2016 Tensorflow: a system for large-scale machine learning *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* pp 265–83
- [50] Semiconductor N 2019 nRF52840 product specification v1.1
- [51] Inc A 2009 Cortex-m4 technical reference manual (available at: <https://developer.arm.com/documentation/ddi0439/b/>)
- [52] Keil Keil mdk version 5: component-based software development (available at: [www2.keil.com/mdk5](http://www2.keil.com/mdk5))
- [53] MathWorks 2020 Deep learning with matlab coder: generate C++ code for deep learning neural networks
- [54] Inc G 2019 Generate model interfaces with tensorflow lite code generator (available at: [www.tensorflow.org/lite/inference\\_with\\_metadata/codegen](http://www.tensorflow.org/lite/inference_with_metadata/codegen))
- [55] Liu X, Zhang M, Subei B, Richardson A G, Lucas T H and Van der Spiegel J 2015 The PennBMBI: design of a general

- purpose wireless brain-machine-brain interface system *IEEE Trans. Biomed. Circuits Syst.* **9** 248–58
- [56] Subasi A and Gursoy M I 2010 EEG signal classification using PCA, ICA, LDA and support vector machines *Expert Syst. Appl.* **37** 8659–66
- [57] Winterhalder M, Maiwald T, Voss H, Aschenbrenner-Scheibe R, Timmer J and Schulze-Bonhage A 2003 The seizure prediction characteristic: a general framework to assess and compare seizure prediction methods *Epilepsy Behav.* **4** 318–25
- [58] Youngerman B E, Khan F A and McKhann G M 2019 Stereoelectroencephalography in epilepsy, cognitive neurophysiology and psychiatric disease: safety, efficacy and place in therapy *Neuropsychiatric Dis. Treat.* **15** 1701–16
- [59] Sze V, Chen Y, Yang T and Emer J S 2017 Efficient processing of deep neural networks: a tutorial and survey *Proc. IEEE* **105** 2295–329
- [60] Taghavi M and Shoaran M 2019 Hardware complexity analysis of deep neural networks and decision tree ensembles for real-time neural data classification *2019 9th Int. IEEE/Conf. on Neural Engineering (NER)* pp 407–10
- [61] Cecotti H and Graser A 2011 Convolutional neural networks for p300 detection with application to brain-computer interfaces *IEEE Trans. Pattern Anal. Mach. Intell.* **33** 433–45
- [62] Kwak N, Muller K and Lee S 2017 A convolutional neural network for steady state visual evoked potential classification under ambulatory environment *PLoS One* **12** e0172578
- [63] Fahimi F, Zhang Z, Goh W B, Lee T S, Ang K K and Guan C 2019 Inter-subject transfer learning with an end-to-end deep convolutional neural network for EEG-based BCI *J. Neural Eng.* **16** 026007
- [64] Tsinalis O, Matthews P M, Guo Y and Zafeiriou S 2016 Automatic sleep stage scoring with single-channel EEG using convolutional neural networks (arXiv:1610.01683)
- [65] Li S, Zhou W, Yuan Q and Liu Y 2013 Seizure prediction using spike rate of intracranial EEG *IEEE Trans. Neural Syst. Rehabil. Eng.* **21** 880–6
- [66] Zhang Z and Parhi K K 2016 Low-complexity seizure prediction from iEEG/sEEG using spectral power and ratios of spectral power *IEEE Trans. Biomed. Circuits Syst.* **10** 693–706
- [67] Chu H, Chung C K, Jeong W and Cho K H 2017 Predicting epileptic seizures from scalp EEG based on attractor state analysis *Comput. Methods Programs Biomed.* **143** 75–87
- [68] Alotaiby T N, Alshebeili S A, Alotaibi F M and Alrshoud S R 2017 Epileptic seizure prediction using CSP and LDA for scalp EEG signals *Comput. Intell. Neurosci.* **2017** 1240323
- [69] Aarabi A and He B 2017 Seizure prediction in patients with focal hippocampal epilepsy *Clin. Neurophysiol.* **128** 1299–307
- [70] Shahbazi M and Aghajan H 2018 A generalizable model for seizure prediction based on deep learning using CNN-LSTM architecture *2018 IEEE Conf. on Signal and Information Processing (GlobalSIP)* (IEEE) pp 469–73
- [71] Affes A, Mdhaffar A, Triki C, Jmaiel M and Freisleben B 2019 A convolutional gated recurrent neural network for epileptic seizure prediction *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* vol 11862 LNCS pp 85–96