Early Seizure Detection with an Energy-Efficient Convolutional Neural Network on an Implantable Microcontroller

Maria Hügle*, Simon Heller†, Manuel Watter*, Manuel Blum*, Farrokh Manzouri‡, Matthias Dümpelmann‡
Andreas Schulze-Bonhage‡, Peter Woias†, Joschka Boedecker*
*Dept. of Computer Science, Faculty of Engineering, University of Freiburg
{hueglem, watterm, mblum, jboedeck}@informatik.uni-freiburg.de
†Dept. of Microsystems Engineering, Faculty of Engineering, University of Freiburg
{simon.heller, woias}@imtek.uni-freiburg.de
‡Epilepsy Center, Medical Center, Faculty of Medicine, University of Freiburg
{farrokh.manzouri, matthias.duempelmann, andreas.schulze-bonhage}@uniklinik-freiburg.de

Abstract—**Implantable, closed-loop devices for automated early detection and stimulation of epileptic seizures are promising treatment options for patients with severe epilepsy that cannot be treated with traditional means. Most approaches for early seizure detection in the literature are, however, not optimized for implementation on ultra-low power microcontrollers required for long-term implantation. In this paper we present a convolutional neural network for the early detection of seizures from intracranial EEG signals, designed specifically for this purpose. In addition, we investigate approximations to comply with hardware limits while preserving accuracy. We compare our approach to three previously proposed convolutional neural networks and a feature-based SVM classifier with respect to detection accuracy, latency and computational needs. Evaluation is based on a comprehensive database with long-term EEG recordings. The proposed method outperforms the other detectors with a median sensitivity of 0.96, false detection rate of 10.1 per hour and median detection delay of 3.7 seconds, while being the only approach suited to be realized on a low power microcontroller due to its parsimonious use of computational and memory resources.**

Index Terms—**Electroencephalography, Epilepsy, Seizure Detection, Responsive Neurostimulation, Convolutional Neural Network, Low Power Microcontroller**

I. INTRODUCTION

Epilepsy is one of the most common neurological diseases [1] and a high percentage of patients with epilepsies are refractory to pharmaceutical therapy [2]. A new treatment option for patients with intractable epilepsy is closed-loop brain stimulation [3], with an additional advantage of only short intermittent interventions compared to traditional continuous pharmaceutical therapy. In order to interrupt seizures, a seizure detection algorithm can be used to trigger the intervention using intracranial Electroencephalography (EEG) data. Research on automatic seizure detection started with the objective of reducing workload for the review of long-term recordings in epilepsy monitoring units and moved together with the development of seizure prediction algorithms [4] towards the application in implantable devices [5], [6]. Most of the approaches are based on handcrafted feature selection [7], [8] or rules designed by experts [9]. Driven by the success of deep learning, more recent approaches use deep convolutional networks or recurrent networks [10]–[12]. However, most of these architectures are too demanding for the implementation on an implantable hardware platform. There are some approaches which are using small convolutional neural networks with only few layers and a small number of weights. The EEGNet of Lawhern et al. [13] is in principle transferable to an implantable hardware. However, the hardware implementation itself was not addressed in their work. Kiral-Kornek et al. [14] proposed two architectures for seizure prediction. One convolutional neural network, which was evaluated on a comprehensive data set and another, which can be implemented on a TrueNorth neuromorphic chip from IBM.

In this paper, we propose SeizureNet, which uses efficient layer combinations and has state-of-the-art detection performance. SeizureNet bridges the gap to an implant for seizure detection based on deep learning. To the best of our knowledge, we have designed the first convolutional neural network for seizure detection specifically for an implantable ultra-low power microcontroller. The proposed architecture exhibits low runtime and memory usage, but maintains high sensitivity in combination with a low false positive rate and a short detection delay for a successful stimulation in a later closed-loop application.

After describing the hardware and the dataset in the two following sections, we define the seizure detection problem and explain our detection pipeline, including preprocessing, model architecture, training and performance evaluation in section IV. We then show results in section V for the actual seizure detection performance, as well as a comparison of hardware properties such as runtime, memory, and energy consumption of our model and four other baselines. We also discuss limitations of the seizure detection device, before concluding in section VI.

II. HARDWARE

For the hardware implementation of the network, a low power microcontroller MSP430FR5994 from Texas Instruments is used. Due to its power consumption of $118\,\mu$A/MHz in active mode and $0.5\,\mu$A in standby mode, it is suitable for the application in an implantable device where a heating of the surrounding tissue must be avoided. A further great advantage of the MSP430FR series is its ferromagnetic nonvolatile memory (FRAM). With a low-power consumption and fast write speed, a swift storage of hidden layer activations of a neural network can be implemented. However, the FRAM also limits the maximum clock speed of the controller as its reading speed is limited to $8\,$MHz. It is possible to run the controller with higher clock speeds but only with additional wait states for the CPU leading to a lower power efficiency. Another useful feature for the implementation of convolution layers is the 32-bit hardware multiplier of the controller, enabling power efficient multiply and accumulate (MAC) operations without CPU intervention.

III. DATASET

The dataset used is the *Epilepsiae* database [15], containing long-term continuous intracranial EEG data. We evaluate our approach on 24 patients. Each recording has a duration between five and eleven days and contains the measurement of approximately 100 intracranial and scalp electrodes originally sampled with or resampled to $f_s = 256\,$Hz. During the two weeks, the evaluated patients had between $6 - 92$ seizures. To limit the amount of data for our experiments, we consider 100 minutes segments of the recordings around the seizures.

For every patient, we consider a subset of $E = 4$ electrodes, which are selected a priori by expert epileptologists to cover the seizure onset zone(s). In case that less than four electrodes display the initial ictal EEG pattern, neighboring channels are included for seizure detection. The total number of electrodes is limited due to hardware limitations.

IV. METHODS

A. Seizure Detection Problem

Seizure detection can be modeled as time-series classification, where we classify ictal phases (seizures) and interictal phases (non-seizures). To create the inputs for the convolutional network, we process sliding windows over the EEG data $D \in \mathbb{R}^{T \cdot E}$, where T is the recording duration and E is the number of electrodes. 1 second samples are created as input features $x_t \in R^I$ with input dimensionality $I = \lfloor T/S \rfloor \cdot E \cdot f_s$ at time point t, where f_s is the sampling frequency and S the stride. The corresponding labels are $y \in \{0,1\}^{\lfloor T/S \rfloor}$, where $y_i = 1$ indicates a seizure at the end of sample i and $y_i = 0$ an interictal sample. The window lengths are chosen to keep the runtime of a forward pass low. To train our models, an overlap of 50% is used, which equals a stride of $S = 0.5$ seconds. Due to the hardware runtime limitations, we evaluate our models with a stride of $1\,$s.

B. Preprocessing

Compared to the conventional scalp EEG, intracranial EEG is less prone to artifacts like the pick-up of the electrocardiogram or electromyogram. However, a careful removal of noise and drifts in the data can facilitate the subsequent pattern classification task. Multiple preprocessing steps are performed to remove signal components which are not carrying relevant information and to adapt the signal statistics to intra-individual fluctuations. The power line noise at $50\,$Hz is removed by applying a notch-filter on the raw EEG data to exclude this frequency. Subsequently, a highpass filter ($\geq 0.1\,$Hz) removes slow drifts. The data is then rescaled by dividing through the rolling 10 minutes standard deviation to account for non-stationarity in the source. Preliminary experiments showed this setting to perform well. The rolling mean and standard deviation can be computed for a time point t and window size N as:

$$\mu_t^N(x_{t-N:t}) = \frac{1}{N} \sum_{i=t-N}^{t} x_i$$

$$\sigma_t^N(x_{t-N:t}) = \frac{1}{N} \sum_{i=t-N}^{t} \left(x_i - \mu_t^N(x_{t-N:t}) \right)^2$$

The input data is then normalized as follows:

$$\hat{x}_t = \tanh \left(0.2 \cdot \frac{x_t}{\sigma_t^{N_w}(x_{t-N_w:t})} \right)$$

where $N_w = 153600$, which equals 10 minutes of data. We normalize the scaled and standardized data via the hyperbolic tangent to reduce the influence of outliers and artifacts while maintaining a quasi-linear relation for most of the input distribution.

C. Model Architecture

In order to find a good model architecture, we evaluated the runtime and memory requirements for various layer types like convolutions, dense layers, pooling layers and activations. The architecture of SeizureNet is shown in Table I. The proposed network is a deep convolutional network with alternating convolutional and pooling layers.

Lawhern et al. [13] proposed convolutions over electrodes in the first layer. They used kernels of size $(E, 1)$, which is similar to approaches such as common spatial patterns [16]. We extend this by convolving over the electrodes and time, so that we can learn spatio-temporal patterns efficiently in one layer. In the last layer, we use 1×1 convolutions instead of a fully-connected layer. This was introduced in [17] and is a parameter-efficient way to reduce dimensions [18]. Rectified Linear Units (ReLu) are used in all hidden layers. Further, batch normalization is used after the convolutions. During training, dropout regularization is applied to reduce overfitting.

D. Training

To deal with the high imbalance of ictal and interictal samples, we use an oversampling technique. Mini-batches are created by randomly picking ictal samples with probability p,

and interictal samples otherwise. In order to learn to detect seizures as early as possible, we weight the ictal samples in the loss function according to their distances to the seizure onset. The weights decrease linearly from 1.0 for the onset to 0.0 for the seizure offset.

In all experiments, the patient-specific models are trained. For evaluation, we use 3-fold cross validation. Each model is trained with a batch size of 256 for $1.5 \cdot 10^5$ steps (batches of samples), with a sampling probability $p = 0.1$ for the number of seizures in each batch. For optimization, we use Adam [19] with a learning rate of 10^{-3} and the binary cross-entropy loss.

E. Detection Performance Evaluation

It is non-trivial to evaluate a seizure detection system. Mainly, three objectives should be optimized:

- The *sensitivity* is defined as the ratio of actually detected seizures to the total number of seizures.
- The *detection delay* is calculated as the mean delay over all detected seizures. For each detected seizure, the delay is defined as the expired time between the electrographic seizure onset identified through visual inspection by a domain expert, and the first algorithm-based detection of the seizure.
- The *false positive rate* is the number of false detections per hour (*fp/h*).

F. Approximations

For our hardware implementation, we have to approximate the rolling 10 minutes standard deviation due to memory limitations. We first approximate the rolling 10 minutes mean $\mu_t^{N_w}$ by computing a grand mean $\tilde{\mu}_t^{N_w}$ over 1-second means $\bar{\mu}_k$. Doing so, only $K = 600$ means have to be stored for the targeted duration, while introducing a negligible error by ignoring the most recent samples until they form a new 1

second mean $\bar{\mu}_1$. The approximated mean can then be used to compute the 10 minutes standard deviation $\tilde{\sigma}_t^{N_w}$:

$$\bar{\mu}_k = \mu_{t-k \cdot f_s}^{f_s} \left(x_{t-k \cdot f_s : t-(k-1)f_s} \right)$$

$$\tilde{\mu}_t^{N_w} = \mu_t^K \left(\bar{\mu}_{1:K} \right)$$

$$\left(\tilde{\sigma}_t^{N_w} \right)^2 = \frac{1}{N_w} \sum_{i=t-N}^{t} \left(x_i - \tilde{\mu}_t^{N_w} \right)^2$$

with $k \in \{1, \ldots, K\}$.

Further, we use a linear approximation of the hyperbolic tangent in the preprocessing to avoid the need for a lookup-table:

$$\text{lintanh}(x) := \begin{cases} x/1.2 & -1.2 \leq x \leq 1.2 \\ 1 & x > 1.2 \\ -1 & x < -1.2 \end{cases}$$

G. Comparison to other Approaches

For seizure detection, it is hard to compare approaches without evaluating on the same framework, due to factors like different evaluation methods, different datasets, omitted patients and patient-specific or across-patients training. Hence, we reimplemented all our baselines. We compare the performance of our architecture to three other convolutional neural networks, which we find most similar to our approach. In order to evaluate fairly, we use the same preprocessing and the sampling method for all convolutional networks. The architectures of the network baselines are shown in Table II. Further, we compare to the performance of a support vector machine approach from the literature using handcrafted features [8].

1) EEGNet: The first baseline is EEGNet, a small convolutional network by Lawhern et al. [13]. The EEGNet shows robust performance across four different brain-computer interface classification tasks. It consists of three convolutional layers and two max pooling layers, where the first convolution

TABLE I
ARCHITECTURE OF SEIZURENET FOR E INPUT ELECTRODES.

Layer	Operation	Output
	Input ($E \times 256$)	$E \times 256 \times 1$
1	$20 \times$ Conv2D ($E \times 17$)	$1 \times 240 \times 20$
2	MaxPool2D (1×4)	$1 \times 60 \times 20$
	Dropout (0.2)	$1 \times 60 \times 20$
3	$10 \times$ Conv2D (1×5)	$1 \times 56 \times 10$
4	MaxPool2D (1×4)	$1 \times 14 \times 10$
	Dropout (0.2)	$1 \times 14 \times 10$
5	$10 \times$ Conv2D (1×5)	$1 \times 10 \times 10$
6	MaxPool2D (1×2)	$1 \times 5 \times 10$
	Dropout (0.2)	$1 \times 5 \times 10$
7	$10 \times$ Conv2D (1×5)	$1 \times 1 \times 10$
	Dropout (0.2)	$1 \times 1 \times 10$
8	$1 \times$ Conv2D (1×1)	$1 \times 1 \times 1$
9	Sigmoid	1
	Total Number of Parameters: 3,621	

TABLE II
ARCHITECTURE OF THE BASELINE NETWORKS.

EEGNet	Kiral-Kornek et al.	Acharya et al.
Input ($E \times 256$)	Input ($32 \times 32 \times E$)	Input ($E \times 4097$)
$16 \times$ Conv2D ($E \times 1$)	$16 \times$ Conv2D (3×3)	$4 \times$ Conv2D ($E \times 6$)
Transpose	MaxPool2D (2×2)	MaxPool2D (1×2)
Dropout (0.25)	Dropout (0.7)	$4 \times$ Conv2D (1×5)
$4 \times$ Conv2D (2×32)	$32 \times$ Conv2D (3×3)	MaxPool2D (1×2)
MaxPool2D (2×4)	MaxPool2D (2×2)	$10 \times$ Conv2D (1×4)
Dropout (0.25)	Dropout (0.7)	MaxPool2D (1×2)
$4 \times$ Conv2D (8×4)	$32 \times$ Conv2D (3×3)	$10 \times$ Conv2D (1×4)
MaxPool2D (2×4)	MaxPool2D (2×2)	MaxPool2D (1×2)
Dropout (0.25)	Dropout (0.7)	$15 \times$ Conv2D (1×4)
$1 \times$ Conv2D (1×1)	Dense (32)	MaxPool2D (1×2)
GlobalMaxPool2D	Dropout(0.5)	Dense (50)
		Dense (20)
Sigmoid Output Layer		
Total Number of Parameters:		
957	15,665	96,220

estimates a set of spatial filters over the electrodes. To adapt their approach to our framework, we replace the softmax regression output layer by a sigmoid activation.

2) Kiral-Kornek et al.: In [14], a convolutional network is evaluated for patient-specific seizure prediction based on spectrograms. To compare with their architecture, we use their approach for seizure detection instead of prediction. They proposed a network consisting of three alternating convolutional layers and three max-pooling layers using 32×32 spectrograms as input. As activation function, they use an Exponential Linear Unit (ELU) [20]. To provide a proof-of-concept that they can implement their seizure prediction system on a low-power system, they adapt their network architecture to a 18-layer binary neural network, consisting only of convolutional layers and dropout. This architecture can run on the IBM *TrueNorth Neurosynaptic System* chip [21]. However, because their network uses binary weights, they need more layers and thus more parameters to reach the same precision. In total, they use over 4.2 million parameters, which would cost more than $525\,\mathrm{kB}$ for binary weights. Because this already exceeds our memory limit, we only use their small architecture as a baseline. To adapt this approach to seizure detection, we use 1 second windows instead of 30 second windows and Short-Time Fourier Transformation (STFT) to generate the spectrograms. Further, we focus on the comparison of the network architectures and thus don't include time of day as an input feature to keep the input consistent across methods.

3) Acharya et al.: In [11], a 13-layer convolutional neural network is trained across patients on one electrode to classify interictal, ictal, and preictal phases (phases directly before a seizure). They trained their model on a subset of the *Epilepsiae* database, consisting of 5 epileptic patients and additionally 5 healthy subjects. The architecture takes an input window of length 4097, which equals to $16\,\mathrm{s}$ for a sampling rate of $256\,\mathrm{Hz}$. As activation function, they use a Leaky ReLu. To stick to our framework, we adapt their network so that it uses 4 electrodes as input by changing the first convolutional layer from a 6×1 convolution to 6×4. Further, we use a sigmoid output layer. For training, an overlap of 50% is used. As for the other approaches, the performance is evaluated with a 1 second stride.

4) SVM: Besides deep learning approaches, we compare our method to the support vector machine (SVM) proposed in [8]. For feature extraction, they use the EMD algorithm [22], which is a signal processing method for nonlinear and non-stationary time series. The SVM uses the variance of Intrinsic Mode Functions (IMFs) as input features and a Radial Basis Function (RBF) as kernel. A post processing step is applied, which classifies a seizure only if the SVM has classified a number of consecutive epochs as epileptic activity. In the original study, eight consecutive epochs were required for detecting a seizure. Since this would lead to a high detection delay for online detection steps of one second, we skip this step. Further, they use only 20 seconds of ictal and 100 seconds of interictal data for training. In contrast to that, we use all available seizures of the training set and downsample the remaining interictal data to the ratio proposed in the original paper. In our experiments, we use the first three extracted IMFs as features for the SVM.

V. RESULTS

In order to evaluate our approach, we first compare the detection performance for all architectures. Further, the ROC AUC score is computed to show the discriminative properties of the models. Since the AUC score is independent of the classifier threshold, it is a suitable additional metric to compare the different architectures. However, since seizures only have to be detected once but as fast as possible, the per-sample sensitivity and thus the AUC score of the classifier is not as important as the other metrics. After that, we show how the performance metrics can be influenced by varying the classification threshold. This offers the possibility of manual (re-)adjustment of the detection device during a study with real patients. Then, we investigate the effect of the preprocessing approximations. Finally, we evaluate all network-based models regarding their memory and speed for our hardware implementation. For this purpose, we analyze the requirements for the preprocessing and the successive layers of the networks.


Fig. 1. Detection performances over all patients for a classifier threshold of 0.5. Median delay and false positive rate are shown in logarithmic scale.

The detection performance of our model and all baselines is shown in Fig. 1 for a classifier threshold of 0.5. To show the overall performance, we calculate the metrics separately for each patient and summarize the distribution in the figure. SeizureNet has robust and high sensitivities for all patients. With a median sensitivity of 0.96, it outperforms all other

architecture except for the network of Kiral-Kornek et al., which has a median sensitivity of 0.98. Their network, however, shows extremely high false positives rates, with outliers up to 464 false positives per hour. Because of the large window size of $16\,\mathrm{s}$, the lowest median fp/h of 0.04 is achieved by the network of Acharya et al. However, this window leads to a high median delay of $24.3\,\mathrm{s}$ (unacceptable for *early* seizure detection applications) and a low sensitivity of 0.67. The second best false positive rate is achieved by EEGNet due to its low sensitivity. The SVM has a good median sensitivity, but a higher delay and a higher false positive rate than SeizureNet. With the best median AUC score of 0.89 and a good balance between false positives and delay, SeizureNet shows the best and most robust detection performance.

Fig. 2 shows the influence of the classifier threshold on the performance metrics. Higher thresholds correlate with reduced sensitivity and increased delays, but also prove to be less prone to false alarms. The highest median sensitivity of 0.96 is achieved with a classifier threshold 0.5, with a respective median delay of $3.7\,\mathrm{s}$ and 10.1 false positives per hour. A low median fp/h of 0.7 can be achieved with a high classifier threshold of 0.9. However, this threshold decreases the sensitivity to 0.75 and increases the median delay to $10.8\,\mathrm{s}$. Of course, the classifier threshold can also be varied for the other architectures (for the SVM, the proposed postprocessing step can be used instead). However, as indicated by the high AUC score, SeizureNet has the best performance independent of the threshold.


Fig. 2. Influence of classifier threshold on the detection performance of SeizureNet. Median delay and false positive rate are shown in logarithmic scale.

A. Approximations

Fig. 3 illustrates the performance loss due to preprocessing for the proposed approximations for 10 patients. We compare to another approximation, where we replace the rolling mean with 0. The used highpass filter roughly centers the signal around the zero line which allows us to dispense with a mean estimation and the standard deviation can be approximated by

$$\left(\hat{\sigma}_t^{N_w}\right)^2 \approx \frac{1}{N_w} \sum_{i=t-N_w}^{t} x_i^2$$

In our experiments, using the 10 minutes standard deviation and the hyperbolic tangent shows the best results. The best approximation is using the grand mean with a median delay of $3.4\,s$ and a fp/h of 14.1 and an AUC score of 0.94. The
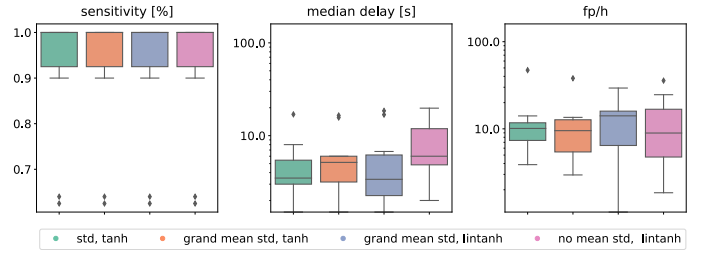

Fig. 3. Detection performance of the proposed preprocessing and various approximations evaluated on a subset of 10 patients.

zero mean approximation offers a big advantage in runtime and memory reduction, however, it increases the median delay considerably to $6\,s$.

B. Hardware Requirement Analysis

The theoretical runtime and memory requirements of all convolutional networks for our hardware implementation are shown in Fig. 4. Besides SeizureNet, only the network of Kiral Kornek et al. would actually be implementable on our microcontroller. However, the preprocessing of this network is extensive due to the STFT and has a 38% higher runtime. The high runtimes of EEGNet and Acharya et al. make an implementation on our realtime detection device impossible. For EEGNet, this is caused by zero-padded convolutions, which are not reducing the dimensionality of the input. The runtime for the network of Acharya et al. is mainly affected by the large input window.

The required memory of the hardware implementation is specified by the number of parameters of the networks, a buffer for the rolling window of the preprocessing, lookup-tables and two buffers, saving the inputs and outputs of the hidden layers. The size of the buffers is determined by the layer with the largest input and output size. Regarding the memory, all networks besides the network proposed of Acharya et al. are theoretically implementable on our device. The limiting factor of this network is the use of fully-connected layers, which require 74% of the total memory.

Besides SeizureNet, none of the networks were actually implemented on hardware platforms, so we cannot compare hardware efficiency and power consumption. Only Kiral-Kornek et al. implemented an adapted network on the IBM *TrueNorth* chip. This chip has a consumption of $40 - 70\,\mathrm{mW}$ to power 1 Mio neurons, resulting in a power consumption of $40 - 70$ nW per neuron. With 107400 neurons, a power consumption of $4.3 - 7.5$ mW can be estimated for the adapted network. For the preprocessing and forward pass of SeizureNet, we measured a power consumption of $850\,\mu W$ which is 5 to 8.8 times lower.

The SVM approach was excluded from this comparison as non-parametric approaches behave differently for each patient. While test-time predictions scale linearly in time and memory with the number of support vectors, the required amount grows rapidly with the number of seizures considered in the training data. Due to the varying computational demands for each patient, it is difficult to find a fair setting for a comparison
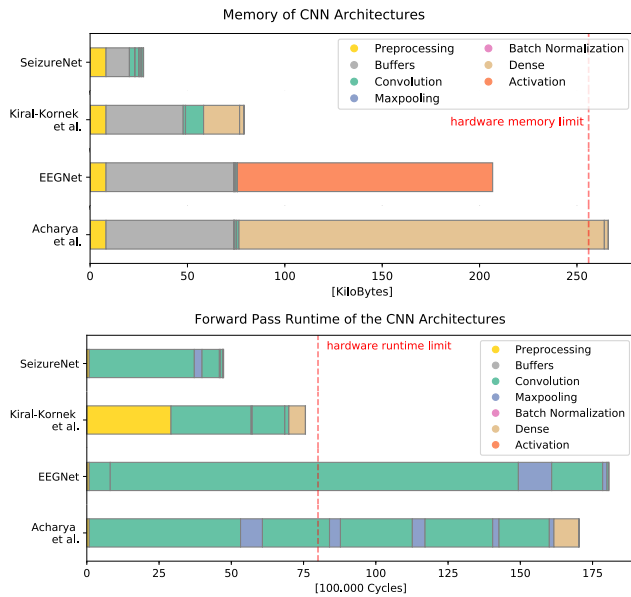
Fig. 4. Memory (top) and runtime (bottom) requirements for SeizureNet and the baseline architectures. Runtime blocks are ordered according to their execution in the forward pass. Layers with few parameters or cycles are not visible.

against the network approaches that have a constant load across patients. Additionally, the preprocessing in [8] would require simplification to be efficiently implementable and change the method substantially.

C. Limitations of the Approach

Tuning the networks to early and sensitive detection of electroencephalographic seizure patterns (Fig. 5) occurs at the cost of higher false positive rates. Several components can have contributed to this: a) selection of short EEG segments of 1 s decreases detection delay but increases chances to falsely classify artefacts as ictal patterns; longer windows of analyses can ameliorate this (note the low FP rate in [11] using a 16 s window); b) subclinical ictal electroencephalographic patterns resembling ictal patterns by definition, but not accompanied by clinical phenomena can be detected which may occur more frequently than clinically manifest seizures [23]. These detections should not be called false detections, and there may even be reasons to use such detections to trigger interventions in a closed-loop device setting. Integration of more EEG channels may allow for an estimation of the probability of clinical correlates of the ictal event.

Expert review of some missed seizure showed epileptic auras with unclear electrographic correlates, which allowed neither clear visual nor algorithmic detection (Fig. 6 bottom). Additional investigations are needed to identify the patterns relevant for seizure detection during the evolution of the electrographic ictal event, which not always coincide with the seizure onset pattern, but rather with alterations of the ictal pattern in the course of recruitment and spread (Fig. 6 top and middle). Remarkably, SeizureNet had also good detection


Fig. 5. Early detected seizure with the repetitive spiking onset pattern: predictions one minute around the seizures and the normalized electrode signal of one electrode.
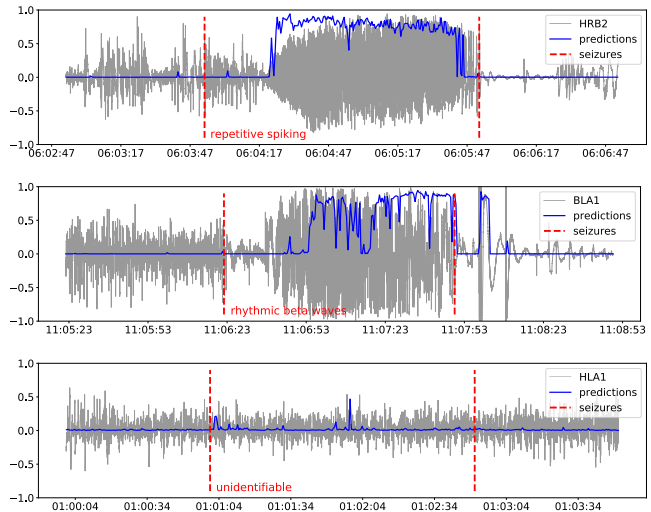

Fig. 6. Late detections (top, middle) and missed seizure (bottom): predictions one minute around the seizures and the normalized electrode signal of one electrode.

latencies for seizures with onset patterns, which are difficult to detect, for example for amplitude depression [24].

VI. CONCLUSION

We have introduced SeizureNet, a convolutional neural network for online seizure detection with state-of-the-art performance. Empirical evaluation of our approach demonstrates its suitability for practical realization on an implantable low-power microcontroller for clinical applications. The considered approximations to preprocessing and architecture choices preserve performance sufficiently while leaving over computational resources for further improvement. Candidates for future work in this direction are the distillation of network ensembles [25] or the improvements in quantized, low-precision neural networks [26].

REFERENCES

[1] J.W. Sander and S.D. Shorvon. Epidemiology of the epilepsies Methodological issues. *Journal of Neurology, Neurosurgery & Psychiatry*, 61(5):433–443, 1996.

[2] Patrick Kwan and Martin J. Brodie. Early identification of refractory epilepsy. *N Engl J Med*, 2000.

[3] Andreas Schulze-Bonhage. Brain stimulation as a neuromodulatory epilepsy therapy. *Seizure*, 548:239–247, nov 2016.

[4] Florian Mormann, Ralph G. Andrzejak, Christian E. Elger, and Klaus Lehnertz. Seizure prediction: The long and winding road. *Brain*, 130(2):314–333, 2007.

[5] Sriram Ramgopal, Sigride Thome-Souza, Michele Jackson, Navah Ester Kadish, Iván Sánchez Fernández, Jacquelyn Klehm, William Bosl, Claus Reinsberger, Steven Schachter, and Tobias Loddenkemper. Seizure detection, seizure prediction, and closed-loop warning systems in epilepsy. *Epilepsy and Behavior*, 37:291–307, 2014.

[6] Mark J. Cook, Terence J. O'Brien, Samuel F. Berkovic, Michael Murphy, Andrew Morokoff, Gavin Fabinyi, Wendyl D'Souza, Raju Yerra, John Archer, Lucas Litewka, Sean Hosking, Paul Lightfoot, Vanessa Ruedebusch, W. Douglas Sheffield, David Snyder, Kent Leyde, and David Himes. Prediction of seizure likelihood with a long-term, implanted seizure advisory system in patients with drug-resistant epilepsy: A first-in-man study. *The Lancet Neurology*, 12(6):563–571, 2013.

[7] Lojini Logesparan, Alexander J. Casson, and Esther Rodriguez-Villegas. Optimal features for online seizure detection. *Medical and Biological Engineering and Computing*, 50(7):659–669, 2012.

[8] Yuxin Zheng, Junming Zhu, Yu Qi, Xiaoxiang Zheng, and Jianmin Zhang. An automatic patient-specific seizure onset detection method using intracranial electroencephalography. 18, 09 2014.

[9] Jean Gotman. Automatic recognition of epileptic seizures in the EEG. *Electroencephalography and clinical Neurophysiology*, pages 530–540, 1982.

[10] Nhan Duy Truong, Anh Duy Nguyen, Levin Kuhlmann, Mohammad Reza Bonyadi, Jiawei Yang, and Omid Kavehei. A generalised seizure prediction with convolutional neural networks for intracranial and scalp electroencephalogram data analysis. *CoRR*, abs/1707.01976, 2017.

[11] U. Rajendra Acharya, Shu Lih Oh, Yuki Hagiwara, Jen Hong Tan, and Hojjat Adeli. Deep convolutional neural network for the automated detection and diagnosis of seizure using eeg signals. *Computers in Biology and Medicine*, 2017.

[12] Pierre Thodoroff, Joelle Pineau, and Andrew Lim. Learning robust features using deep learning for automatic seizure detection. *CoRR*, abs/1608.00220, 2016.

[13] Vernon J. Lawhern, Amelia J. Solon, Nicholas R. Waytowich, Stephen M. Gordon, Chou P. Hung, and Brent J. Lance. Eegnet: A compact convolutional network for eeg-based brain-computer interfaces. *CoRR*, abs/1611.08024, 2016.

[14] Isabell Kiral-Kornek, Subhrajit Roy, Ewan Nurse, Benjamin Mashford, Philippa Karoly, Thomas Carroll, Daniel Payne, Susmita Saha, Steven Baldassano, Terence O'Brien, David Grayden, Mark Cook, Dean Freestone, and Stefan Harrer. Epileptic seizure prediction using big data and deep learning: Toward a mobile system. *EBioMedicine*, 2017.

[15] Juliane Klatt, Hinnerk Feldwisch-Drentrup, Matthias Ihle, Vincent Navarro, Markus Neufang, Cesar Teixeira, Claude Adam, Mario Valderrama, Catalina Alvarado-Rojas, Adrien Witon, Michel Le Van Quyen, Francisco Sales, Antonio Dourado, Jens Timmer, Andreas Schulze-Bonhage, and Bjoern Schelter. The EPILEPSIAE database: An extensive electroencephalography database of epilepsy patients. *Epilepsia*, 53(9):1669–1676, 2012.

[16] B. Blankertz, R. Tomioka, S. Lemm, M. Kawanabe, and K. r. Muller. Optimizing spatial filters for robust eeg single-trial analysis. *IEEE Signal Processing Magazine*, 25(1):41–56, 2008.

[17] Min Lin, Qiang Chen, and Shuicheng Yan. Network in network. *CoRR*, abs/1312.4400, 2013.

[18] J.T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller. Striving for simplicity: The all convolutional net. In *ICLR (workshop track)*, 2015.

[19] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.

[20] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *CoRR*, abs/1511.07289, 2015.

[21] Steven K. Esser, Paul A. Merolla, John V. Arthur, Andrew S. Cassidy, Rathinakumar Appuswamy, Alexander Andreopoulos, David J. Berg, Jeffrey L. McKinstry, Timothy Melano, Davis R. Barch, Carmelo di Nolfo, Pallab Datta, Arnon Amir, Brian Taba, Myron D. Flickner, and Dharmendra S. Modha. Convolutional networks for fast, energy-efficient neuromorphic computing. *CoRR*, abs/1603.08270, 2016.

[22] Norden E. Huang, Zheng Shen, Steven R. Long, Manli C. Wu, Hsing H. Shih, Quanan Zheng, Nai-Chyuan Yen, Chi Chao Tung, and Henry H. Liu. The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 454(1971):903–995, 1998.

[23] Hinnerk Feldwisch-Drentrup, Matthias Ihle, Michel le van Quyen, Cesar Teixeira, Antonio Dourado, Jens Timmer, Francisco Sales, Vincent Navarro, Andreas Schulze-Bonhage, and Bjoern Schelter. Anticipating the unobserved: Prediction of subclinical seizures. *Epilepsy and Behavior*, 22(SUPPL. 1):S119–S126, 2011.

[24] Ralph Meier, Heike Dittrich, Andreas Schulze-Bonhage, and Ad Aertsen. Detecting epileptic seizures in long-term human EEG: a new approach to automatic online and real-time detection and classification of polymorphic seizure patterns. *Journal of clinical neurophysiology : official publication of the American Electroencephalographic Society*, 25(3):119–131, 2008.

[25] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015.

[26] Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental network quantization: Towards lossless cnns with low-precision weights. 2017.