

An Accurate EEGNet-based Motor-Imagery Brain–Computer Interface for Low-Power Edge Computing

Xiaying Wang^{*†}, Michael Hersche^{*†}, Batuhan Tömekce[†], Burak Kaya[†], Michele Magno^{†‡}, Luca Benini^{†‡}

[†]ETH Zürich, Dept. EE & IT, Switzerland

[‡]University of Bologna, DEI, Italy

Abstract—This paper presents an accurate and robust embedded motor-imagery brain–computer interface (MI-BCI). The proposed novel model, based on EEGNet [1], matches the requirements of memory footprint and computational resources of low-power microcontroller units (MCUs), such as the ARM Cortex-M family. Furthermore, the paper presents a set of methods, including temporal downsampling, channel selection, and narrowing of the classification window, to further scale down the model to relax memory requirements with negligible accuracy degradation. Experimental results on the Physionet EEG Motor Movement/Imagery Dataset show that standard EEGNet achieves 82.43%, 75.07%, and 65.07% classification accuracy on 2-, 3-, and 4-class MI tasks in global validation, outperforming the state-of-the-art (SoA) convolutional neural network (CNN) by 2.05%, 5.25%, and 6.49%. Our novel method further scales down the standard EEGNet at a negligible accuracy loss of 0.31% with 7.6× memory footprint reduction and a small accuracy loss of 2.51% with 15× reduction. The scaled models are deployed on a commercial Cortex-M4F MCU taking 101 ms and consuming 4.28 mJ per inference for operating the smallest model, and on a Cortex-M7 with 44 ms and 18.1 mJ per inference for the medium-sized model, enabling a fully autonomous, wearable, and accurate low-power BCI.

Index Terms—Brain–computer interface, motor-imagery, CNN, embedded systems, edge computing

I. INTRODUCTION

Brain–computer interfaces (BCIs) aim to provide a communication and control channel based on the recognition of the subject’s intentions, e.g., when performing motor-imagery (MI), from neural activity typically recorded by noninvasive electroencephalogram (EEG) electrodes [2]. MI-BCI systems are designed to find patterns in the EEG signals and match the signal to the motor motion that was imagined by the subject. Such information could enable communication for severely paralyzed users, control of a wheelchair [3], or assistance in stroke rehabilitation [4].

MI-BCIs are still susceptible to errors mostly due to high inter- and intra-subject variance in EEG data [5], [6], resulting in low classification accuracy. Traditional methods approach this challenge with robust feature extractors, typically filter bank common spatial pattern (FBCSP) [7] or Riemannian covariances [8], and classify these features with linear discriminant analysis (LDA) or support vector machines (SVMs) [6]. Convolutional neural networks (CNNs) have been proposed as a competitive solution in EEG classification, while requiring fewer parameters to learn and being computationally cheaper in inference than traditional BCI methods [9], [1]. However,

^{*}X. Wang and M. Hersche contributed equally to this work as first authors. Corresponding emails: {xiaywang, herschmi}@iis.ee.ethz.ch

today’s CNN models are designed to be executed on a CPU or GPU, requiring EEG data to be transmitted from the sensor node to an external compute engine through wired or wireless communication. Due to their computational complexity and resource requirements, those models have been predominantly confined to cloud computing with high-performance computers rather than used in real-world BCI applications, where latency, privacy, and wearability are crucial requirements besides the accuracy [10], [11]. Recently, a new generation of wearable BCIs is attracting the academic and industrial researchers. An increasing number of battery-operated wearable solutions, using microcontroller units (MCUs), are proposed to bring computing capabilities towards the “edge” to perform real-time near-sensor computation [10], [12], [13], [14]. Edge computing and near-sensor computation offer the following advantages: 1) lower energy consumption for the data transmission between sensors and remote processing; 2) longer battery lifetime; 3) significantly shorter latency compared to remote computation; 4) user comfort; 5) security and privacy improvements, as the data are processed locally and only little information is transmitted wirelessly if necessary.

On the other hand, edge computing poses several challenges when it needs to match the requirement of long-term battery operation, mandatory in wearable devices, and to continuously perform complex BCI models (e.g., CNNs) with a low-power processor. For instance, the ARM Cortex-M series is the most popular family of low-power processors used in embedded wearable devices [15]. Those MCUs allow several hours lifetime with a small-scale battery, but they have a resource-constrained architecture. For example, an ARM Cortex-M4F processor offers few KB of RAM and million operations per second (MOPS) in a power envelope of few mW [16]. The more recent ARM Cortex-M7 provides an even better performance up to 300-400 MOPS with a higher power consumption of few hundred mW. To achieve the goal of deploying complex and accurate CNN models on these tiny microprocessors, the models need to be re-thought and redesigned with the above-mentioned constraints in mind. Moreover, many researchers have demonstrated for computer vision that reducing the model size with clever network optimization techniques does not always cause a performance degradation [15], [17].

This paper proposes a novel embedded model for MI-BCI that focuses on bringing the next generation of edge BCI on autonomous wearable systems. The main contributions of the paper are as follows:

- We propose a novel embedded MI-BCI model which

outperforms the state-of-the-art (SoA) model on the Physionet EEG Motor Movement/Imagery Dataset [18]. The model is based on EEGNet architecture [1] and achieves a global validation accuracy of 82.43%, 75.07%, and 65.07% on 2-, 3-, and 4-class MI task, which is 2.05%, 5.25%, and 6.49% higher than the SoA CNN [19].

- We further propose methods to reduce the memory footprint for the execution of the model by temporal downsampling, channel reduction, and narrowing down the time window considered for performing one classification, without significant loss in accuracy. This allows us to target low-power embedded devices with very tight constraints.
- We evaluate experimentally the benefits of our model in terms of energy consumption, latency, and accuracy on two different platforms: ARM Cortex-M4F and Cortex-M7. We compare the two platforms executing the inference of 4-class MI with accurate measurements.

To the best of our knowledge, no previous work has evaluated MI-BCI on these low-power MCUs using CNNs by considering both runtime and power measurements besides the classification accuracy. Finally, we release open-source code developed in this work¹.

II. RELATED WORK

The recent literature on MI-BCIs is very rich, mostly considering feature extraction and classifiers separately. EEG signals are typically pre-processed using spectral and spatial filters followed by log-energy feature calculation, better known as FBCSP [7], [20]. The multi-spectral features are classified using either LDA, regularized LDA, or SVMs [6].

Alternatively, the feature extractor and classifier can be combined and trained simultaneously with a CNN. Today, CNNs are among the most accurate BCI architectures and demonstrated impressive performance [19], [9], [1]. Schirrmeyer et al. [9] provide an elaborate study on CNN architectures for MI-BCI, where the small Shallow ConvNet achieves an accuracy of 73.59% on the 4-class MI-BCI competition IV-2a [5]. With its temporal and spatial filters followed by square-log activation, Shallow ConvNet can be interpreted as a tunable variant of FBCSP.

Due to limited amount of data provided in the MI-BCI competition IV-2a dataset containing recordings of only 9 subjects with 144 MI-trials per class, a variant of Shallow ConvNet has been trained and validated in [19] on the much larger Physionet EEG Motor Movement/Imagery Dataset [18], with recordings of 109 subjects with 21 MI-trials per class which is overall $\approx 2\times$ the amount of MI-trials. The model consists of around 305 thousand trainable parameters and is trained and validated globally in 5-fold cross-validation (CV) across subjects. It has achieved SoA accuracy of 80.38%, 69.82%, and 58.58% on 2-, 3-, and 4-class MI on that dataset. Additionally, the global models are adjusted for every subject using subject-specific transfer learning (SS-TL), which further improved the accuracy by 6.11%, 9.42%, and 9.93%. The main differences to the original Shallow ConvNet are the use of

ReLU instead of square-log activation and the splitting of the final classification layer into two fully connected layers, which increases the number of trainable parameters.

Another smaller, yet robust, CNN architecture is EEGNet [1], which achieved the same accuracy as the winner of the BCI competition IV-2a on 4-class MI [5]. The main difference to Shallow ConvNet is that EEGNet uses fewer feature maps, spatial separable convolutions, and more pooling layers, which reduces the number of weights as well as the feature map sizes. Its flexibility and small size, however, comes at the cost of significantly lower accuracy, e.g., 67% for 4-class MI. Efforts have been made to modify EEGNet by changing the pooling layers and expanding the network achieving 72% accuracy with subject-specific models [21].

Most of these models are evaluated remotely offline, without considering the possibility to bring the computation closer to the sensors, where the data is acquired. Few studies have shown embedded implementations using traditional MI-BCIs with separate feature extractors [22], [23], [24], but, to the best of our knowledge, no previous work has demonstrated accurate embedded MI-BCI on low-power MCUs using CNNs, which offers better accuracy at lower latency. In this paper, we propose a CNN model based on EEGNet to perform MI classification on Physionet dataset [18]. Our proposed model improves the 4-class MI accuracy by 6.49% in global validation while requiring two orders of magnitude less parameters and reducing the memory footprint during inference by a factor of $4.6\times$ compared to the current SoA on this dataset. In order to target resource-constrained low-power embedded systems, we further study model reduction methods to test the limitations of the proposed model architecture in terms of the sampling frequency, the number of EEG channels, and the length of the input signals. We implement two reduced models that are within the resource constraints of two popular low-power MCUs with ARM Cortex-M4F and M7 and measured the runtime and energy consumption.

III. DATASET DESCRIPTION

We use the publicly available Physionet EEG Motor Movement/Imagery Dataset [18] containing EEG recordings of 109 subjects. Four subjects are discarded due to variability in the number of trials, resulting in 105 subjects to be finally used. The EEG signals were recorded with the BCI2000 system [25] using 64 channels sampled at 160 Hz. The subjects performed motor movement and MI tasks, however, in this study we solely focus on the classification of MI. Every subject participated in three runs for MI of left fist (L) against right fist (R), and three runs for MI of both fists (B) against both feet (F). One run lasts 120 s and consists of 14 MI trials according to the timing scheme shown in Fig. 1. This results in 21 trials per class per subject. A baseline run provides resting data (0), where the subjects did not receive any cues for 60 s while having their eyes open. In order to get trials with resting data, we extract windows of 3 s from the baseline run. As done in [19], we distinguish between 2-, 3-, and 4-class MI using L/R, L/R/0, and L/R/0/F MI tasks, respectively.

¹<https://github.com/MHersche/eegnet-based-embedded-bci>




Fig. 1. Trial paradigm [19] of Physionet EEG Motor Movement/Imagery Dataset.

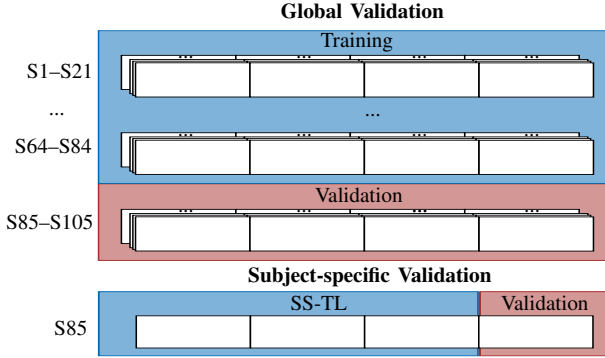


Fig. 2. Training and validation on Physionet EEG Motor Movement/Imagery Dataset. Global validation is done via 5-fold CV over the subjects, whereas subject-specific validation additionally includes transfer learning (SS-TL) in 4-fold validation on the corresponding subject, e.g., on S85.

A. Validation methodology

Fig. 2 illustrates the validation methodology inspired by [19], which distinguishes global from subject-specific validation. The global validation accuracy is determined by 5-fold CV across the subjects, i.e., training on 4/5 of the subjects and validating on the remaining, unseen 1/5 of the subjects. In SS-TL, the global model is further adjusted by doing transfer-learning on part of the subject’s data and validated on the remainder. This validation is done with 4-fold CV on every subject.

In the example of Fig. 2, a global model is first trained on subjects S1–S84 and validated on S85–S105, yielding the first fold accuracy of the global validation. SS-TL is then applied for S85–S105 on every subject separately in 4-fold CV, always starting with the global model from S1–S84.

IV. METHODS

This section introduces the novel embedded MI-BCI model proposed in this paper that matches the memory and complexity constraints of low-power MCUs with high accuracy. We first describe how the compact EEGNet [1] is applied and evaluated on the Physionet Dataset [18], and propose methods to further reduce the memory requirements of EEGNet.

A. EEGNet on Physionet Dataset

Fig. 3 shows the architecture of EEGNet for 4-class MI on the Physionet Dataset. The first 3 s of the MI cue are used for classification, which is the interval [0 s, 3 s] according to Fig. 1. The input feature map represents the EEG signal in time domain with $N_s=480$ samples ($3\text{ s}\times 160\text{ Hz}$) and $N_{ch}=64$ EEG channels. The samples are filtered in the time domain using 1-D convolutions, in the spatial domain

with depthwise convolutions, in the time-spatial domain with separable convolutions, and finally classified with a fully connected layer. Furthermore, EEGNet uses exponential linear unit (ELU) activation and average pooling in the time domain.

Table I gives further insights into the architecture of EEGNet. Here, the number of input samples, EEG channels, kernel size of the first average pooling are kept variable; they have a direct influence on the number of parameters as well as feature map sizes. The last two columns show that EEGNet is indeed very compact: it requires to learn only 3,204 weights, which is two orders of magnitude less than the SoA CNN [19]. However, large feature maps need to be stored during operation. Assuming we need to be able to store at least two consecutive feature maps at any time, the maximum number of stored features is the sum of the input and first layer, i.e., $N_s \times N_{ch} + N_s \times N_{ch} \times 8 = 276,480$ features.

Training and evaluation of the EEGNet are implemented using Keras with TensorFlow (version 1.11) backend. The model is trained with Adam optimizer for 100 epochs with batch size of 16 and a fixed learning rate scheduler, setting the learning rate to 0.01, 0.001, and 0.0001 at epochs 0, 20, and 50, respectively. In SS-TL, the model is trained for 5 more epochs. All training hyperparameters were determined via 5-fold CV on the training set of the first fold of the global validation set (i.e., S1–S84) for 4-class MI, and kept the same for 2- and 3-class MI as well as all reduced EEGNet configurations.

B. Embedded implementation

For the evaluation on embedded processors, we choose two MCUs from STMicroelectronics: B-L475-IOT01A2 with an ARM Cortex-M4F processor at 80 MHz with 128 KB of SRAM and 1 MB of Flash memory and STM32F756 Nucleo-144 with an ARM Cortex-M7 processor at 216 MHz with 320 KB SRAM and 1 MB of Flash. Both MCUs utilize digital signal processors and floating-point units. We then use the X-CUBE-AI expansion package of STM32CubeMX [26] to deploy the trained models on the selected MCUs.

Based on Table I and considering 32-bit floating-point numbers, the estimated Flash memory needed for storing the parameters of the model is around 13 KB, whereas the RAM requirement for the to largest consecutive feature maps is roughly 1.05 MB. With these configurations, the model can not be executed on the selected low-power MCUs. As mentioned in the previous subsection, the output of the first layer requires most of the memory. To overcome this limitation, we reduce the input data size by:

- 1) downsampling the EEG data in the time domain. MI activities cause brain oscillations mainly within the μ (8–14 Hz) and β (14–30 Hz) bands [9]. Some CNNs have shown to learn temporal filters that cover the gamma (71–91 Hz) band, however, the main information was still extracted from μ and β oscillations [19]. Therefore, we downsample the signals by a factor of $ds=2$ or $ds=3$, which restricts us to maximal oscillations of 40 Hz or 26 Hz, respectively. The temporal filter and the pooling kernel sizes are scaled to $N_f = \lceil 128/ds \rceil$ and $N_p = \lceil 8/ds \rceil$. This way, the network is expected to




Fig. 3. EEGNet [1] in standard configuration for 4-class MI on the Physionet Motor Movement/Imagery dataset. A window of 3 s ($N_s = 480$ samples) with $N_{ch} = 64$ channels is classified at the time.

TABLE I

Detailed description of EEGNet in MI classification. N_s is the number of input samples in time domain, N_{ch} the number of EEG channels, N_{cl} the number of classes, N_f the filter size of first temporal filter, and N_p the pooling length. For each map, n is the number of filters, p the padding strategy, k the kernel size, and s the stride. The last two columns show number of parameters and feature map size of standard configuration, i.e. $N_s = 480$, $N_{ch} = 64$, $N_{cl} = 4$, $N_f = 128$, $N_p = 8$.

Layer	Type	n	p	k	s	Parameters	Output Shape	Standard Configuration		
								Parameters	Feature Map Size	
ϕ^1	Conv2d	8	same	$N_f \times 1$	1×1	$8N_f$	$N_s \times N_{ch} \times 8$	1,024	245,760	
	BatchNorm2d	-	-	-	-	32		32		
ϕ^2	DepthConv2d	16	valid	$1 \times N_{ch}$	1×0	$N_{ch} \cdot 16$	$1 \times N_s // N_p \times 16$	1,024	960	
	BatchNorm2d	-	-	-	-	64		64		
	EluAct	-	-	-	-	-		-		-
	AvgPool2d	-	valid	$N_p \times 1$	$N_p \times 1$	-		-		-
ϕ^3	SepConv2d	16	same	16×1	1×1	512	$1 \times N_s // N_p // 8 \times 16$	512	120	
	BatchNorm2d	-	-	-	-	64		64		
	EluAct	-	-	-	-	-		-		-
	AvgPool2d	-	valid	8×1	1×8	-		-		-
ϕ^4	FC	4	-	-	-	$(N_s // N_p // 8 \cdot 16 + 1)N_{cl}$	N_{cl}	484	4	
	SoftMaxAct	-	-	-	-	-		-		-
Total (inkl. Input Feature Map)						$672 + 16N_{ch} + 8N_f + (2N_s/N_p + 1)N_{cl}$	$N_s(9N_{ch} + 18/N_p) + N_{cl}$	3,204	277,564	




Fig. 4. Electrode configurations.

learn similarly to the original model after the depthwise convolution, independent on the downsampling factor.

- 2) using a subset of electrode channels. The BCI2000 system conforms to the 10-10 international system electrode placement with $N_{ch}=64$ electrodes. We reduce the number of electrodes to $N_{ch}=19$ by taking the widely used 10-20 international system electrode placement, from which we exclude A1 and A2. As an intermediate configuration, we add the channels to cover the whole

region of the brain equally reaching $N_{ch}=38$ electrodes. We also investigate the case with only 8 electrodes based on the EEG headset by Bitbrain. Fig. 4 shows the 64-, 38-, 19-, and 8-electrodes configurations.

- 3) decreasing the time window of the signal used for each classification. We reduce the input signal from 3 s to 2 s or 1 s after the start of the MI cue (ref. Fig. 1). Noteworthy, this approach reduces the delay of the system in addition to the model size reduction.

We study the impact on classification accuracy of each reduction approach, testing different configurations to choose the best combination in terms of accuracy and memory footprint for further deployment on both selected MCUs.

V. EXPERIMENTAL RESULTS

This section assesses the proposed methods on the Physionet Motor Movement/Imagery Dataset. We measure the classification accuracy as the ratio between correct classified trials over the total number of trials.

A. Global vs. Subject-specific MI Classification

Table II compares the average classification accuracy of the global and subject-specific model of EEGNet with the baseline CNN proposed in [19]. In global validation, EEGNet outperforms the baseline CNN by 2.05%, 5.25%, and 6.49% on 2-, 3-, and 4-class MI, respectively. EEGNet does not improve as significantly as the baseline CNN when applying SS-TL: the accuracy increases by 1.89%, 5.00% and 5.76% on EEGNet and by 6.11%, 9.43%, and 9.93% on the baseline

TABLE II

Classification accuracy (%) on Physionet EEG Motor Movement/Imagery Dataset using global model (global) or subject-specific model with transfer learning (SS-TL). The top accuracies for each configuration are marked bold.

	Dose et al. [19]		This work	
	global	SS-TL	global	SS-TL
2 classes	80.38	86.49	82.43	84.32
3 classes	69.82	79.25	75.07	80.07
4 classes	58.58	68.51	65.07	70.83

CNN. Due to already high accuracy of EEGNet in global validation, however, the accuracy in subject-specific validation is still 0.82% and 2.32% higher than the baseline CNN in 3- and 4-class MI and only 2.17% lower in 2-class MI.

B. EEGNet Model Reduction

Table III studies the impact on the classification accuracy in 2-, 3-, and 4-class MI on global validation when reducing EEGNet by temporal downsampling, channel reduction, and narrowing the classified time window. Only one reduction approach is applied at a time; the remaining configurations are kept according to the standard EEGNet.

As expected, downsampling has a negligible effect on the accuracy with a maximum decrease among all MI tasks of 0.32% and 1.25% at downsampling factor $ds=2$ and $ds=3$, respectively. Even though part of the β band is ignored at $ds=3$ due to the cut-off frequency of the anti-aliasing filter at ≈ 26 Hz, the accuracy does not drop significantly. This result confirms that the significant information in this dataset is contained in α and lower β bands for the MI task. When reducing the number of EEG channels to $N_{ch}=38$, the accuracy decreases only marginally by a maximum of 0.95%. However, further reduction to $N_{ch}=19$ and $N_{ch}=8$ significantly affects the performance with a maximum accuracy decrease of 2.66% and 6.52%, respectively. Similar trends can be seen when narrowing down the time window used to do one classification: the accuracy decreases by a maximum of 1.62% with a temporal window of $T=2$ s, and by 3.6% with $T=1$ s. As already mentioned in the previous section, narrowing the time window brings additional advantages in shorter classification delays, and thus, provides a trade-off between accuracy and delay to be chosen by the user.

Next, we test all combinations of reduction methods in order to find the best configuration in terms of accuracy vs. memory footprint. Fig. 5 shows the global 4-class accuracy of all reduction combinations, excluding the $N_{ch}=8$ configuration due to the large drop in accuracy. We consider only the memory footprint required to store input and first layer features with 32-bit floating-point representation, since the number of features is two orders of magnitudes higher than the parameters, as pointed out in Table I. As delay might be an additional constraint for model selection, the configurations are marked according to the time window used for classification. EEGNet outperforms the baseline CNN in most configurations: it has at least $4.6\times$ lower memory footprint (i.e., <1.05 MB vs. 4.80 MB), while achieving higher classification accuracy in most cases. We select two EEGNet configurations on the pareto-optimal curve, which satisfy the




Fig. 5. Global accuracy on 4-class MI vs. RAM requirements for storing feature maps of reduced configurations of EEGNet. The chosen configurations with highest accuracy while staying below the practical RAM limits are “Model 1” (blue circle) with $ds=3$, $N_{ch}=38$, and $T=1$ s for M4, and “Model 2” (cyan circle) with $ds=3$, $N_{ch}=38$, and $T=2$ s for M7.

memory constraints of the chosen Cortex-M4F and M7 MCUs. Both configurations use a downsampling of $ds=3$ and $N_{ch}=38$ channels; they only differ in the time window choosing $T=1$ s at 62.51% accuracy and 72 KB RAM requirements for Cortex-M4, and $T=2$ s at 64.76% with 143 KB for Cortex-M7. This corresponds to an accuracy loss of 2.51% at $15\times$ model reduction when operating EEGNet on the M4, and 0.31% loss at $7.6\times$ model reduction on the M7, compared to the standard EEGNet configuration. We name them “Model 1” and “Model 2”, respectively, in shorthand.

C. MCU Implementation

We deploy the selected models using STM32CubeMX v5.3 with X-CUBE-AI 5.0.0 package extension on STM32L475VG B-L475E-IOT01A with an ARM Cortex-M4F processor and STM32F756ZG Nucleo-144 with an ARM Cortex-M7 processor and measure the power consumption with a Keysight N6705C power analyzer. In order to have optimal performance, we enable the core instruction and data caches and ART accelerator sub-system to speed up instruction fetch accesses. We deploy Model 1 on both Cortex-M4F and Cortex-M7 for comparison, and Model 2 only on Cortex-M7 due to memory constraints.

The Cortex-M7 offers the highest performance in ARM Cortex-M processor family. In fact, as can be seen in Table IV, it takes around 6 cycles per multiply-and-accumulate (MACC) operation, which is around $1.8\times$ faster than the Cortex-M4F. However, the power consumption is $3.1\times$ higher than Cortex-M4F at the same frequency of 80 MHz. Cortex-M7 can run up to 216 MHz, which reduces the latency of Model 1 by a factor of $4.9\times$ compared to Cortex-M4F at a price of almost $2\times$ more energy consumption. Model 2 has the lowest accuracy loss (i.e., 0.31%) after model reduction, but it can fit only into the Cortex-M7 processor. Running at the highest frequency, it takes around 44 ms and 18.1 mJ per inference.

TABLE III

Classification accuracy (%) using global validation. The standard EEGNet (Fig. 3) is reduced either by downsampling in time domain, reducing the number of channels, or narrowing the time window for a single classification.

	Standard	Downsampling		Channel reduction			Time window	
		$ds=2$	$ds=3$	$N_{ch}=38$	$N_{ch}=19$	$N_{ch}=8$	$T=2$ s	$T=1$ s
2 classes	82.43	82.11	81.97	81.86	81.95	78.07	81.11	79.86
3 classes	75.07	74.78	73.82	74.12	72.41	68.99	73.45	71.47
4 classes	65.07	64.81	64.77	64.65	62.55	58.55	64.13	63.51

TABLE IV

ARM Cortex-M4F vs. M7 on 4-class MI. Both Model 1 (62.51% accuracy) and Model 2 (64.76%) use a downsampling of $ds=3$ and $N_{ch}=38$ channels, the former has time window $T=1$ s, while the latter $T=2$ s.

	Model 1		Model 2	
	M4F	M7	M7	M7
	@80 MHz	@80 MHz	@216 MHz	@216 MHz
MACC	761,956	1,509,220		
ROM size [KB]	6.61	7.12		
RAM size [KB]	70.27	139.20		
Cycles/MACC	10.59	5.77	5.78	6.27
Power [mW] @ 3.3V	42.44	131.41	412.76	413.06
T/inference [ms]	100.84	54.99	20.40	43.81
En./inference [mJ]	4.28	7.23	8.42	18.1

VI. CONCLUSION

We propose an embedded model based on EEGNet for low-power MI-BCIs. The proposed model achieves 2.05%, 5.25%, and 6.49% higher accuracy than the SoA CNN on 2-, 3-, and 4-class MI classification, while requiring two orders of magnitude less memory for storing model parameters and $4.6\times$ less memory for feature maps during inference execution. We reduce the input feature map by downsampling in the temporal and spatial domain as well as narrowing down the time window and relax the memory requirements by $7.6\times$ at 0.31% accuracy loss, and by $15\times$ at 2.51% loss. We demonstrate the performance of the proposed models on two commercial MCUs. In particular, the implemented models execute in around 44 ms consuming 18.1 mJ per inference on an ARM Cortex-M7 and in 101 ms using 4.28 mJ on an ARM Cortex-M4F processor, making them suitable for a battery-operated real-time wearable system to continuously perform online MI classification.

ACKNOWLEDGMENT

This project was supported in part by the Swiss Data Science Center PhD Fellowship under grant ID P18-04 and in part by ETH Research Grant 09 18-2.

REFERENCES

- [1] V. J. Lawhern, A. J. Solon *et al.*, "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," *Journal of Neural Engineering*, vol. 15, no. 5, p. 056013, 2018.
- [2] B. Graimann, B. Allison *et al.*, "Brain-Computer Interfaces: A Gentle Introduction." Springer, Berlin, Heidelberg, 2009, pp. 1–27.
- [3] Y. Yu, Z. Zhou *et al.*, "Self-Paced Operation of a Wheelchair Based on a Hybrid Brain-Computer Interface Combining Motor Imagery and P300 Potential," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 12, pp. 2516–2526, 12 2017.
- [4] A. A. Frolov, O. Mokienko *et al.*, "Post-stroke Rehabilitation Training with a Motor-Imagery-Based Brain-Computer Interface (BCI)-Controlled Hand Exoskeleton: A Randomized Controlled Multicenter Trial." *Frontiers in neuroscience*, vol. 11, p. 400, 2017.
- [5] M. Tangermann, K.-R. Müller *et al.*, "Review of the BCI Competition IV." *Frontiers in neuroscience*, vol. 6, p. 55, 2012.
- [6] F. Lotte, L. Bougrain *et al.*, "A review of classification algorithms for EEG-based brain-computer interfaces: a 10 year update," *Journal of Neural Engineering*, vol. 15, no. 3, p. 031005, 2018.
- [7] K. K. Ang, Z. Y. Chin *et al.*, "Filter Bank Common Spatial Pattern (FBCSP) in brain-computer interface," *Proceedings of the International Joint Conference on Neural Networks*, pp. 2390–2397, 2008.
- [8] F. Yger, M. Berar *et al.*, "Riemannian Approaches in Brain-Computer Interfaces: A Review," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 10, pp. 1753–1762, 10 2017.
- [9] R. T. Schirmer, J. T. Springenberg *et al.*, "Deep learning with convolutional neural networks for EEG decoding and visualization," *Human Brain Mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [10] J. Chen and X. Ran, "Deep Learning With Edge Computing: A Review," *Proceedings of the IEEE*, vol. 107, no. 8, 2019.
- [11] L. Angrisani, P. Arpaia *et al.*, "A Single-Channel SSVEP-Based Instrument With Off-the-Shelf Components for Trainingless Brain-Computer Interfaces," *IEEE Transactions on Instrumentation and Measurement*, vol. 68, no. 10, pp. 3616–3625, 2019.
- [12] M. Guermandi, S. Benatti *et al.*, "A Wearable Device for Minimally-Invasive Behind-the-Ear EEG and Evoked Potentials," in *Proc. IEEE BioCAS*, Oct. 2018, pp. 1–4.
- [13] V. Kartsch, G. Tagliavini *et al.*, "BioWolf: A sub-10mW 8-channel Advanced Brain Computer Interface Platform with a 9-core processor and BLE connectivity," *IEEE Transactions on Biomedical Engineering*, 2019.
- [14] X. Wang, M. Magno *et al.*, "FANN-on-MCU: An Open-Source Toolkit for Energy-Efficient Neural Network Inference at the Edge of the Internet of Things," *IEEE Internet of Things Journal*, 2020.
- [15] L. Lai, N. Suda *et al.*, "Cmsis-nn: Efficient neural network kernels for arm cortex-m cpus," *arXiv:1801.06601*, 2018.
- [16] M. Eggmann, S. Mach *et al.*, "A RISC-V Based Open Hardware Platform for Always-On Wearable Smart Sensing," in *Proc. IEEE 8th IWASI*, 2019, pp. 169–174.
- [17] F. N. Iandola, S. Han *et al.*, "SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size," *arXiv:1602.07360*, 2016.
- [18] A. L. Goldberger, L. A. N. Amaral *et al.*, "PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals," *circulation*, vol. 101, no. 23, pp. e215–e220, 2000.
- [19] H. Dose, J. S. Møller *et al.*, "An end-to-end deep learning approach to MI-EEG signal classification for BCIs," *Expert Systems with Applications*, 2018.
- [20] M. Hersche, T. Rellstab *et al.*, "Fast and Accurate Multiclass Inference for MI-BCIs Using Large Multiscale Temporal and Spectral Features," in *Proc. IEEE 26th EUSIPCO*, 2018, pp. 1690–1694.
- [21] A. Uran, C. van Gemeren *et al.*, "Applying Transfer Learning To Deep Learned Models For EEG Analysis," *arXiv:1907.01332*, 2019.
- [22] C. M. McCrimmon, J. L. Fu *et al.*, "Performance Assessment of a Custom, Portable, and Low-Cost Brain-Computer Interface Platform," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 10, pp. 2313–2320, 10 2017.
- [23] K. A. Condori, E. C. Urquiza *et al.*, "Embedded Brain Machine Interface based on motor imagery paradigm to control prosthetic hand," in *2016 IEEE ANDESCON*. IEEE, 10 2016, pp. 1–4.
- [24] K. Belwafi, O. Romain *et al.*, "An embedded implementation based on adaptive filter bank for brain-computer interface systems," *Journal of Neuroscience Methods*, 2018.
- [25] G. Schalk, D. McFarland *et al.*, "BCI2000: A General-Purpose Brain-Computer Interface (BCI) System," *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 6, pp. 1034–1043, 6 2004.
- [26] STMicroelectronics, "STMCube.AI," in *STMCubeMX*, 2019.