

Towards a Domain-Specific Neural Network Approach for EEG Bad Channel Detection

V. Kumaravel^{1,2}, F. Paissan¹ and E. Farella¹

1. E3DA Unit, Fondazione Bruno Kessler, Trento, Italy

2. CIMeC, University of Trento, Italy

{vkumaravel, fpaissan, efarella}@fbk.eu

Electroencephalogram (EEG) is prone to several artifacts that often lead to misclassification of neural features in Brain-Computer Interfaces (BCI) [1]. In recent years, deep learning (DL) techniques have been successfully used to decode brain activities using EEG data [2]. Most of the proposed methods use a hand-made preprocessing strategy to remove EEG noise before feeding the data into Neural Networks (NNs) [2, 3]. In contrast, others do not explicitly mention the preprocessing performed.

We believe a dedicated NN-based preprocessing pipeline integrated with the downstream analysis pipeline might significantly improve the overall system performance. This abstract reports our preliminary results obtained with a 2D Convolutional Neural Network - *cleanEEGNet*, towards this goal. As an initial step, we trained *cleanEEGNet* to detect noisy channels (i.e., electrodes with a low signal-to-noise ratio). We carefully chose the model hyperparameters (such as kernel size and stride) to mimic the conventional detection of bad channels performed via visual inspection. An open-source dataset from OpenNeuro [4] with annotated bad channels is used to train and validate the network. As the proposed approach is the first work based on NN for this task, we chose four state-of-the-art automated conventional methods for comparison: FASTER [5], Clean Raw Data (CRD) [6], HAPPE [7], and HAPPILEE [8]. Finally, we made the source code freely available on GitHub [9].

In EEG, bad channels manifest mainly due to a poor connection between the electrode and the scalp (e.g., electrode pop-out), or due to movement artifacts. The presence of such bad channels negatively impacts the efficiency of overall artifact removal, and propagates noise information to all the channels. While some bad channels are easy to detect upon visual inspection (such as flat-line channels), some others are more difficult to identify. Therefore, researchers rely on statistical features such as power spectral density (PSD), kurtosis, root mean square (RMS) variance, etc., to detect and interpolate/remove bad channels. Figure 1 shows a sample segment of multi-channel EEG data of 5 seconds duration in which bad channels are highlighted in red.

The main obstacle in solving the bad channel detection (BCD) problem using NNs is that the dataset is highly unbalanced (i.e., non-proportional number of positive and negative classes). Thus, learning algorithms need to be tailored to classify both classes well to correct the dataset bias. Unfortunately, synthetically generating EEG signals (i.e., for data augmentation) is not an option without compromising neural information, since for every bad channel in the generated set there are many more good channels. To correct the high bias in the dataset towards good channels, we weighted the loss function such that an error in classifying the less probable outcome is considered worse than an error on the most probable outcome [10]. Moreover, we used a two-factor loss function, combining weighted binary cross-entropy (BCE) and soft F1 score, which is robust to the unbalanced dataset. Our analysis showed that this combination of cost functions improves both balanced accuracy and F1 score with respect to only soft-F1 score or only BCE-based optimization.

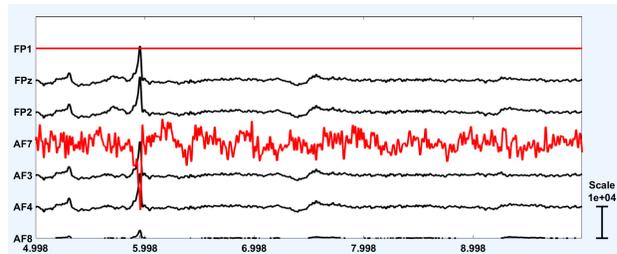


Figure 1. Sample data of length 5 seconds (x-axis) with bad channels 'FP1', and 'AF7' highlighted in red. The channels in black are good channels.

Layer	# filters	Shape	Activation
Input	-	(1, 62, 330000)	
Padding	-	(0, 0, 1, 1)	
Conv2D	8	(3, 256)	ReLU
Depth-wise Conv2D	32	(3, 3)	ReLU
Depth-wise Conv2D	64	(3, 3)	ReLU
Global Average Pooling	-	-	-
Dense		64 to 62	

Table 1. *cleanEEGNet* architecture overview. The table does not show batch normalization, which is performed between every convolutional layer and its activation.

The preliminary architecture we propose to solve BCD, *cleanEEGNet*, comprises of three convolutional blocks and one linear layer, as shown in Table 1. In fact, our architecture resembles the well-known EEGNet [2]. As mentioned above, *cleanEEGNet* is tailored to mimic visual inspection. The temporal kernel of the first layer is set to 256 points, thus having a 0.5 s window with an overlap of 0.25 s with the previous one (stride at 128). The choices on these parameters are based on: 1) knowledge from a domain expert who identifies bad channels via visual inspection, 2) temporal and spatial properties of the EEG signal. For example, the Pearson correlation of channels spatially located close to each other is crucial in detecting a noisy channel. The rationale is that noise components in EEG are usually uncorrelated to components of neural origin; hence correlation can be a reliable metric. As an example, both FASTER and CRD pipelines use inter-channel correlation features in their respective algorithms. Henceforth, we hypothesized that our network would learn to distinguish good and bad channels with the aforementioned hyperparameters.

After the first layer, the other convolutional blocks are composed of depth-wise separable convolutions to extract higher-order correlations without drastically increasing the computational burden (i.e., parameter count). In sum, we perform an average on the time axis - to avoid learning the time distribution of bad segments in the signal - and a linear classifier.

To train and benchmark the different bad channel removal pipelines, we used the data from the study [11]. EEG was recorded with a 64-channel HIamp EEG system (g.tec, Shiedlberg, Austria) at a sampling rate of 512 Hz referenced using the linked ears configuration. The electrodes were positioned in the international 10-10 system.

Fourteen subjects (seven females) with a mean age of 23 years took part in the study that involved recordings on three different days, resulting in 130 EEG files (See [11] for more details related to the experimental protocol). All datasets are made freely available by the authors in OpenNeuro database [4]. However, for this work, we could use 113 files mainly due to technical issues in importing European Data Format (EDF) files. We used 11 subjects (90 files which are 80% of total files) for training and validation, and 3 subjects (23 files, i.e., 20% of total files) for testing.

EEG data were filtered at 40 Hz to remove the 50 Hz line noise interference using the default EEGLAB filter. Subsequently, high-pass filtering with the transition edge [0.25, 0.75] Hz was employed to remove DC drifts using the EEGLAB *clean_drifts* function [6]. Moreover, before feeding data to the neural network, we normalized it by using the average max amplitude among datasets, such that the majority of the values stay in the [-1, 1] range.

Table 2. Performance Comparison

Method	Precision	Recall	F1 Score	bACC
FASTER	0.1667	0.1699	0.1683	57.9
CRD	0.1959	0.2876	0.2331	61.69
HAPPE	0.1892	0.4268	0.2622	66.78
HAPPILEE	0.2968	0.3433	0.3184	65.02
cleanEEGNet	0.52	0.59	0.55	72

It is noteworthy that experts labeled only 415 channels out of 7006 channels as bad, resulting in a 6:94 positive class vs. negative class ratio. Given such imbalanced proportion, we used the following unbiased metrics - balanced accuracy (bACC), Precision, Recall, and F1 score, which are defined as

$$Precision = (TP / ((TP + FP))) \quad (1)$$

$$Recall = (TP / ((TP + FN))) \quad (2)$$

$$F1Score = ((2 * Precision * Recall) / (Precision + Recall)) \quad (3)$$

$$bACC = (1/2) * ((TP / (TP + FN)) + ((TN / (TN + FP)))) \quad (4)$$

where TP, TN, FP, FN indicate the number of true positives, true negatives, false positives, and false negatives respectively.

We compared the classification performance of *cleanEEGNet* with the automated traditional approaches, and we found an improvement of 17% in balanced accuracy and 72% improvement in F1 score, as summarised in Table 2.

Among the conventional methods, the most recent algorithm HAPPILEE performs the best with a maximum F1 Score of 0.32. A possible explanation could be that HAPPILEE integrates both HAPPE and CRD pipelines thereby providing a better estimation of bad channels. Moreover, the parameters were calibrated using a validation set of 19 EEG files.

In sum, this paper presents the preliminary results of our ongoing work on a novel domain-specific NN approach to clean EEG noise. As a first step, we focused on the bad channel detection task. We show that our proposed method, *cleanEEGNet*, has successfully learned the task achieving a maximum F1 Score of 0.55. Even if it is seemingly low, *cleanEEGNet* shows a significant performance improvement compared to the widely used traditional tools. Our ongoing work is to enhance the classification accuracy in detecting bad channels and extend our model to identify bad segments present in the data. For this, we will use the Temple University Artifact Corpus [12] with 310 EEG files annotated for 5 kinds of artifacts. In particular, our focus is on building explainable AI models for denoising EEG by exploiting depth-wise separable convolutions, as in [2]. Finally, we believe having an NN-based preprocessing pipeline could benefit the Deep Learning EEG community and enhance the downstream analysis performance as proven in other domains such as audio signal processing [13].

ACKNOWLEDGEMENTS

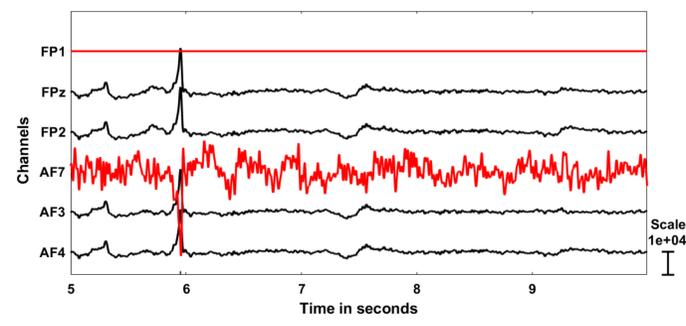
We express our sincere gratitude to the ERC PoC grant NeuroSoNew (842243) that supported this work. We also thank Gianmarco Cerutti and Michele Brugnara for their support and suggestions during the development of this paper.

REFERENCES

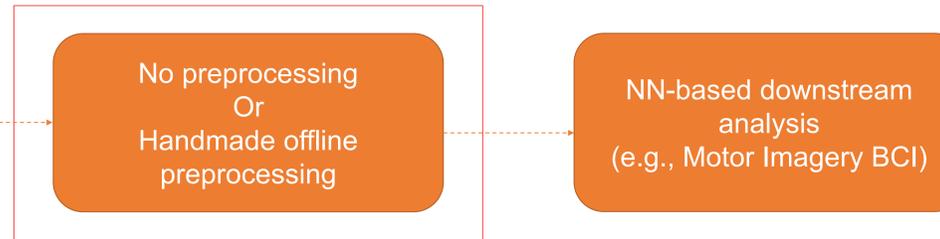
- [1] I. Lazarou et al., "EEG-Based Brain-Computer Interfaces for Communication and Rehabilitation of People with Motor Impairment: A Novel Approach of the 21st Century," in *Journal of Frontiers in Human Neuroscience*, 2018.
- [2] V.J. Lawhern et al., "EEGNet: a compact convolutional neural network for EEG-based brain-computer interfaces," in *Journal of Neural Engineering*, 2018.
- [3] M. Torres et al., "Interpretable SincNet-based Deep Learning for Emotion Recognition from EEG brain activity," in *arXiv preprint arXiv:2107.10790*, 2021.

- [4] C. Schneider et al., “This data was obtained from the Open Neuro database,” in *OpenNeuro*, 2019. [Online]. Available: <https://openneuro.org/datasets/ds002034/versions/1.0.1>
- [5] H. Nolan et al., “FASTER: Fully Automated Statistical Thresholding for EEG artifact Rejection,” in *Journal of Neuroscience Methods*, 2010.
- [6] A. Delorme et al., “EEGLAB Plugin: clean_rawdata, https://github.com/scn/clean_rawdata,” *GitHub*, 2019. [Online]. Available: https://github.com/scn/clean_rawdata.
- [7] L.J. Gabard-Durnam et al., “The Harvard Automated Processing Pipeline for Electroencephalography (HAPPE): Standardized Processing Software for Developmental and High-Artifact Data,” in *Journal of Frontiers in Human Neuroscience*, 2018.
- [8] K.L. Lopez et al., “HAPPILEE: The Harvard Automated Processing Pipeline In Low Electrode Electroencephalography, a standardized software for low density EEG and ERP data,” in *bioRxiv*, 2021.
- [9] V.P. Kumaravel et al., “EEG Bad Channel Detection, https://github.com/vpKumaravel/eeg_bcd/,” *GitHub*, 2021. [Online]. Available: https://github.com/vpKumaravel/eeg_bcd/.
- [10] Q. Qiang et al., “A Nonuniform Weighted Loss Function for Imbalanced Image Classification,” in *Proceedings of the 2018 International Conference on Image and Graphics Processing*, p.78–82, 2018.
- [11] C. Schneider et al., “Real-time EEG Feedback on Alpha Power Lateralization Leads to Behavioral Improvements in a Covert Attention Task,” in *Journal of Brain Topography*, 2019.
- [12] A. Hamid et al., “The Temple University Artifact Corpus: An Annotated Corpus of EEG Artifacts,” in *IEEE Signal Processing in Medicine and Biology Symposium (SPMB)*, 2020.
- [13] M. Ravanelli et al., “Speaker Recognition from Raw Waveform with SincNet,” in *Neural Information Processing Systems (NIPS)*, 2018.

State-of-the-art Scenario



Raw EEG data with bad channels 'FP1', and 'AF7' highlighted in red.



Presence of Bad Channels:

- 1) Bias the learning
- 2) Impacts the performance

Garbage-In Garbage-Out

Proposed Solution: *cleanEEGNet*

Architecture

Layer	# Filters	Shape	Activation
Input	-	(1, 62, 330000)	
Padding	-	(0, 0, 1, 1)	
Conv2D	8	(3, 256)	ReLU
Depth-wise Conv2D	32	(3, 3)	ReLU
Depth-wise Conv2D	64	(3, 3)	ReLU
Global Average Pooling	-	-	-
Dense		64 to 62	

Architecture of *cleanEEGNet*. Batch normalization is performed between every convolution layer and its activation.

Learning

- Bad Channels << Good Channels
- Weighted loss function
- Combining binary cross-entropy (BCE) and soft-F1 score.

$$Perf_{BCE \& \text{soft-F1}} \gg Perf_{BCE} \mid Perf_{\text{soft-F1}}$$

Performance Metrics

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 \text{ Score} = (2 * Precision * Recall) / (Precision + Recall)$$

$$Balanced \text{ Accuracy (bACC)} = \left(\frac{1}{2}\right) * \left(\frac{TP}{TP + FN}\right) + \left(\frac{TN}{TN + FP}\right)$$

Comparison with Traditional Methods

#	Method	Acronym	Features
1	Fully Automated Statistical Thresholding for EEG artifact Rejection	FASTER	Pearson Correlation, Variance, Hurst Exponent
2	EEGLAB's Clean Raw Data	CRD	Pearson Correlation, Variance, Flat-Line Detection
3	The Harvard Automated Processing Pipeline for EEG	HAPPE	Joint probability of mean power in Frequency domain
4	The Harvard Automated Processing Pipeline in Low Electrode EEG	HAPPILEE	HAPPE + CRD

Experimental Data

- Open-source annotated for bad channels
- 14 adult subjects
- 113 EEG files, 62 + 2 REF channels
- Training & Val data – 80% from 11 subjects (90 files)
- Testing data – 20% from 3 subjects (23 files)
- 415 bad out of 7006 channels -> 6:94 ratio

- Band-pass Filtered [1 40] Hz
- Amplitude normalization [-1 1]

Results

