

Predicting Recovery From Coma Following Cardiac Arrest With a Reduced Set of EEG Channels

Nathan T Riek¹, Jonathan Elmer¹, Salah Al-Zaiti¹, Murat Akcakaya¹

¹University of Pittsburgh, Pittsburgh, PA, USA

Abstract

The aim of our work (Univ_Pittsburgh) was to explore the feasibility of using a convolutional neural network (CNN), with a reduced set of EEG channels, fused with a Random Forest to predict coma patient outcomes. This work is part of the ‘Predicting Neurological Recovery from Coma After Cardiac Arrest: The George B. Moody PhysioNet Challenge 2023’.

A 1D-CNN based on the ResNet-18 model was used to detect specific patterns in the EEG unique to either a good or poor outcome for the coma patient. To reduce dimensionality, electrodes were grouped into 5 regions. The CNN was fused with a Random Forest trained on patient features.

The CNN and Random Forest model achieved True Positive Rates (TPRs) of 0.50 +/- 0.09 using 5-fold cross validation within the training set, 0.809 on the training set, 0.448 on the validation set, and 0.530 on the test set. Our team ranked 14th out of 36 teams.

This work demonstrated the feasibility of grouping EEG channels to reduce dimensionality in the prediction of coma recovery. Future work should explore the use of different model architectures with the reduced set of EEG channels to achieve even higher performance.

1. Introduction

It is common for patients who have been resuscitated from cardiac arrest to remain in a coma [1,5-8]. Many of these patients have a poor outcome resulting in either severe neurological impairment or death but others can recover and regain independence in their daily life [7]. Prognostic tests to discriminate between good and poor patient outcomes are imperfect and decisions on life support withdrawal should be delayed more than 72 hours [7]. This provides an opportunity for machine learning to increase prognostic accuracy or reduce the necessary decision time whether to withdraw life support.

EEG has been used to help predict poor outcomes of comatose patients following cardiac arrest [5-9]. EEG is highly dimensional, making it memory intensive and time consuming to interpret. In order to reduce dimensionality,

most works have extracted temporal features [5-9], while others have explored additional spectral and entropy features [5]. Discrete features may capture useful information but may miss some of the temporal significance that is found in continuous EEG [7].

Some works have attempted to capture temporal changes in the EEG by either using raw EEG in a neural network [6] or by taking discrete measurements at multiple time points [7,8]. One group found that aEEG and suppression ratio are very highly correlated spatially and can therefore be averaged across channels to remove noise and allow for greater temporal resolution [7,8]. In other fields using EEG, the EEG signal has been collapsed spatially by grouping the channels into regions [10] or by using techniques like PCA [11].

One work found an improvement in classifying poor outcomes by employing a multi-modal approach, using both continuous EEG data and discrete features [6].

In this work, we propose reducing spatial dimensionality by taking the median of EEG signals within the same brain region. This technique aims to help with both filtering out artifacts and reducing dimensionality while maintaining high temporal resolution. Furthermore, we also leverage the multi-modal approach of using both continuous EEG data and discrete patient features to improve prediction of good or poor patient outcomes.

2. Methods

This work incorporates a hybrid model approach that fuses the predictions from a convolutional neural network (CNN) and a Random Forest. The CNN is trained on continuous EEG data and the Random Forest is trained on patient information.

2.1. Data and Preprocessing

This work is a part of the 2023 PhysioNet Challenge [1-3]. The data used in this work is from the I-CARE EEG database [2], collected from seven hospitals in the U.S. and Europe. There were 1020 total patients in the dataset, each who had return of spontaneous circulation (ROSC) after cardiac arrest but remained comatose. Included in this

work were 19-channel EEG recordings from each hour following ROSC up to 72 hours and several patient-information features. These features were age, sex, time from cardiac arrest to ROSC, if it was an out-of-hospital cardiac arrest (OHCA), if heart rhythm was shockable, and targeted temperature management (TTM) temperature. The labels for the data were Cerebral Performance Category (CPC) scores ranging from one to five with one being the best patient outcome and five being death. The labels were converted to a binary outcome with CPC scores one or two labeled as good and CPC scores three, four, and five labeled as poor.

For any hour containing multiple EEG recordings, the longest recording was kept, and all others were discarded. Starting with the last available hour of EEG data, any EEG channel with a standard deviation of zero was removed. The remaining EEG data was resampled to 128Hz. Next, the EEG data was bandpass filtered from 0.5 to 30Hz using a hamming-window finite impulse response filter [6]. Baseline drift was removed from each channel by subtracting the median amplitude. The channels were grouped into five different regions: (1) left anterior, (2) right anterior, (3) left posterior, (4) right posterior, (5) all. The channels were grouped by taking the median signal of all available channels within the region. Table 1 provides the channels that were included within each region.

Region	Channels
Left Anterior	Fp1, F7, F3, Fz, T3, C3, Cz
Right Anterior	Fp2, F8, F4, Fz, T4, C4, Cz
Left Posterior	O1, T5, P3, Pz, T3, C3, Cz
Right Posterior	O2, T6, P4, Pz, T4, C4, Cz
All	Fp1, Fp2, F7, F8, F3, F4, Fz, T3, T4, T5, T6, C3, C4, Cz, P3, P4, Pz, O1, O2

Table 1. EEG Channel Groupings. The 19-channel EEG data was reduced to 5 channels by taking the median signal within each of these regions.

Once the EEG data was grouped into five channels, the data was segmented into 2500-point windows corresponding to about 19.5s segments. Any absolute-valued windows that contained at least 1250 consecutive points less than $1e-3$ uV in channel five were removed. The EEG data was then normalized by dividing each channel by their absolute maximum amplitude. If the number of kept windows was less than ten, the entire hour of data was skipped, otherwise it was saved. This process was repeated for the training set until either two hours of data were preprocessed or no more EEG data was available. In the validation/test sets up to six hours of data were preprocessed. The data was concatenated across the preprocessed hours so that the dimensions of the data were number of windows x 5 channels x 2500 datapoints. Table 1 provides a summary of all of the patient data that was included for a Random Forest. The sex variable was

converted to three separate binary variables: male, female, other. Any missing patient features were imputed using the mean values across all patients in the training dataset.

Table 1. Training Data Patient Information.

Patient Feature	Mean (\pm std) or Count (%)
Age (years)	61.17 (\pm 15.6)
Sex	Male: 417 (68.7%), Female: 187 (30.8%), Other: 3 (0.5%)
ROSC (minutes)	23.13 (\pm 13.3)
OHCA	True: 566 (93.2%), False: 41 (6.8%)
Shockable Rhythm	True: 575 (94.7%), False: 32 (5.3%)
TTM (degrees Celsius)	33: 448 (73.8%) 36: 61 (10.0%) None: 98 (16.1%)
Outcome	Good: 225 (37.1%) Poor: 382 (62.9%)
CPC	1: 181 (29.8%) 2: 44 (7.2%) 3: 20 (3.3%) 4: 9 (1.5%) 5: 353 (58.2%)

2.2. Model Architecture

The 1D-Convolutional Neural Network (CNN) model is a 1D version of the ResNet-18 architecture [4] and was developed using PyTorch. The expected input dimensions of the model are batch size x 5 channels x 2500 datapoints. The model performs binary classification between good and poor outcomes. The model architecture is shown in Figure 1. Each convolution layer was followed by a batch normalization and ReLU activation function. The ResNet model was selected because of the skip/residual connections, which help avoid the vanishing gradient problem [4]. The cross-entropy loss function and Adam optimizer were selected.

The Random Forest, trained on patient features, was developed using the scikit-learn library in Python. The expected input dimensions of the Random Forest are number of trials x 8 features.

Both the ResNet-18 and Random Forest were tuned to maximize True Positive Rate (TPR) of predicting poor outcomes. TPR was calculated as the number of correctly identified patients with poor outcomes divided by the total number of patients who had poor outcomes. The TPR was maximized while fixing False Positive Rate (FPR) at a maximum value of 0.05. FPR was calculated as the number of patients incorrectly identified as having a poor outcome divided by the total number of patients who had good outcomes.

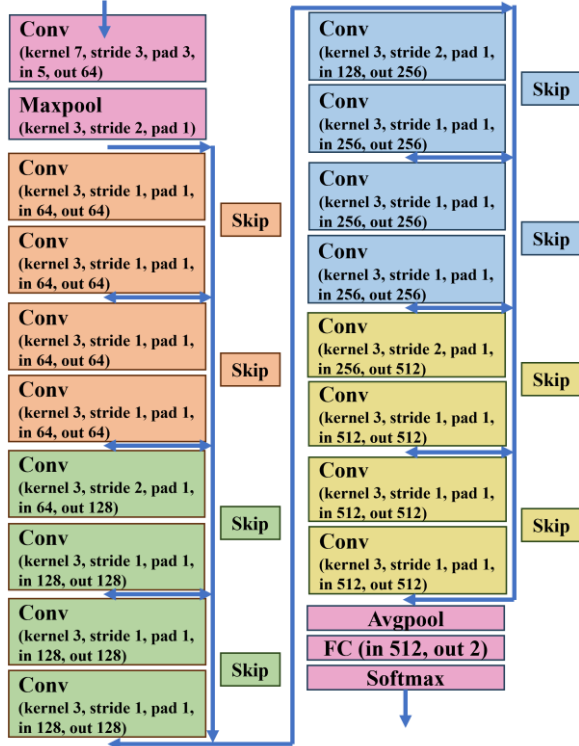


Figure 1. 1D ResNet-18 architecture. The model takes in data of 5 channels x 2500 datapoints. The ResNet-18 begins with a convolution and max-pooling. It then contains 4 different residual blocks colored in orange, green, blue, and yellow. The model ends with an average-pooling, fully connected layer to reduce to 2 classes, and a softmax layer to scale predictions between 0 and 1.

2.2. Model Tuning

The training, validation, and test sets contained data from 607, 107, and 306 patients respectively. Five data folds were created within the training dataset to tune our model architecture. The folds were stratified so that equal numbers of patients with a certain CPC score were in each fold. No two folds contained data from the same patient.

The hyperparameters tuned in the CNN model were batch size and learning rate. A grid search was performed to find the optimal hyperparameters. The optimal hyperparameters were selected based on the highest average TPR from the 5-fold Cross-Validation (CV) on the training set. The batch size was set to 128, and the learning rate was set to 0.01. For each of the folds, the model was trained using two hours of data and adjudicated on the left-out fold using six hours of data. In the left-out fold each patient's EEG data was run on the trained model separately. The output of the Softmax function was averaged across all 2500-datapoint windows of the six hours of EEG data so that each patient had one prediction value. Training was completed in ten epochs for each fold. The TPRs were saved for all ten epochs in each of the five

folds. To find the best generalizable stopping criteria, The TPRs were averaged across folds for each of the ten epochs and the number of epochs was set to the epoch with the highest average TPR. Using this method, the stopping criteria was set to eight epochs. Figure 2 shows the average TPR across the folds for each of the ten epochs. The highest average TPR was 0.47 at epoch eight.

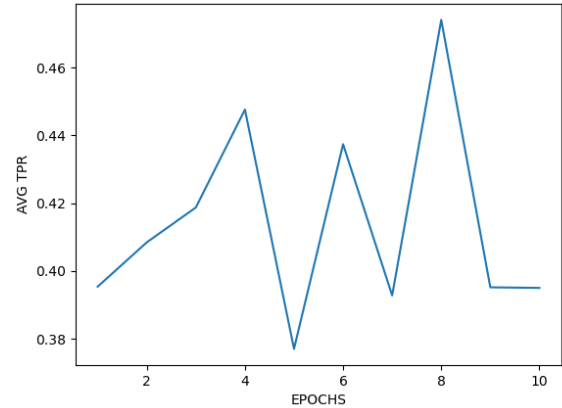


Figure 2. Average ResNet-18 TPR from 5-Fold CV. The model was trained for 10 epochs on each fold, but the eighth epoch had the highest average TPR and so 8 epochs was used for training on the entire training set to be evaluated on the validation and test sets.

The hyperparameters tuned in the Random Forest were the number of trees and the maximum depth of each tree. A grid search was performed using the training set and 5-fold CV yielding the highest average TPR. The number of trees was set to 100 and the maximum depth of each tree was set to 10. Each patient had one set of features, so no averaging of the prediction values was necessary as it was in the CNN.

The prediction values were fused between the CNN and Random Forest using weighted averaging. A grid search was used to find the optimal weights to assign to each of the model's outputs such that the highest average TPR was achieved across the five folds. The best weights were 0.83 for the CNN and 0.17 for the Random Forest. Once all hyperparameters were selected for the CNN, Random Forest, and the fusing of the two models, the model was trained on the entire training set and evaluated on the validation and test sets. Figure 3 shows how the ResNet-18 and Random Forest were fused together.

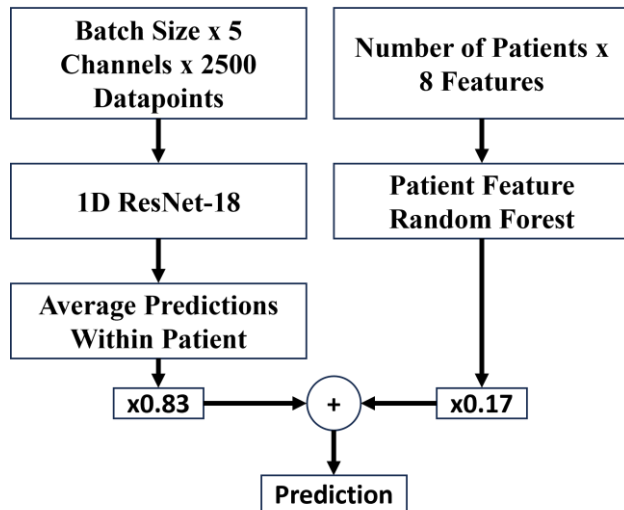


Figure 3. Model Fusion Schematic. The ResNet-18 was trained on EEG data and the Random Forest was trained on patient features. The predictions of the two models are combined using a weighted fusion, resulting in one prediction per patient.

3. Results

The fused ResNet-18 and Random Forest model was evaluated on the training set using 5-fold CV, and then was evaluated on the training set, hidden validation set, and hidden test set. These scores were determined after 72 hours from ROSC. For all results, the FPR was fixed at less than 0.05. Within the training set, the model achieved an average TPR of 0.50 ± 0.09 . The trained model on the training set, validation set, achieved TPRs of 0.809, 0.448, and 0.530 respectively. Table 2 organizes these results.

Training CV	Training	Validation	Test	Ranking
0.50 ± 0.09	0.809	0.448	0.530	14/36

Table 2. Model Performance. The results provide TPR at FPR of 0.05. The Test TPR is the official challenge score.

4. Discussion

By fusing a CNN trained on continuous data with a Random Forest trained on discrete data, the model achieved reasonable performance. This model demonstrates that collapsing EEG into brain regions can help with dimensionality reduction when dealing with large datasets. Reducing from 19 to 5 channels allowed us to train on more than three times the temporal information with the set hardware constraints. Collapsing EEG into brain regions can also help filter noise when one channel becomes disconnected or has large artifacts. Future work should be done with collapsed EEG regions.

Fusing models with continuous data and discrete

features is a method that may benefit the prediction of good or poor outcomes in comatose patients following cardiac arrest. This fusion method may also benefit work on other biosignals and should be explored further.

References

- [1] M. A. Reyna, E. Amorim, et al., "Predicting Neurological Recovery from Coma After Cardiac Arrest: The George B. Moody PhysioNet Challenge 2023," *Computing in Cardiology*, vol. 50, pp. 1-4, 2023.
- [2] E. Amorim, W. L. Zheng, M. M. Ghassemi, M. Aghaeval, P. Kandhare, V. Karukonda, J. W. Lee, S. T. Herman, S. Adithya, N. Gaspard, J. Hofmeijer, M. J. A. M. van Putten, R. Sameni, M. A. Reyna, G. D. Clifford, M. B. Westover. The International Cardiac Arrest Research (I-CARE) Consortium Electroencephalography Database. *Critical Care Medicine* 2023 (in press); doi:10.1097/CCM.0000000000006074.
- [3] A. L. Goldberger, et al., "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. 215-220, Jun, 2000.
- [4] K. He, et al., "Deep Residual Learning for Image Recognition," *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770-778, 2016.
- [5] E. Amorim, et al., "Quantitative EEG Reactivity and Machine Learning for Prognostication in Hypoxic-Ischemic Brain Injury," *Clinical Neurophysiology*, vol. 130, no. 10, pp. 1908-1916, Oct, 2019.
- [6] W. L. Zheng, et al., "Predicting Neurological Outcome in Comatose Patients After Cardiac Arrest With Multiscale Deep Neural Networks," *Resuscitation*, vol. 169, pp. 86-94, Dec, 2021.
- [7] J. Elmer, et al., "A Novel Methodological Framework for Multimodality, Trajectory Model-Based Prognostication," *Resuscitation*, vol. 137, pp. 197-204, Apr, 2019.
- [8] J. Elmer, et al., "Group-Based Trajectory Modeling of Suppression Ratio After Cardiac Arrest," *Neurocritical Care*, vol. 25, pp. 415-423, Dec, 2016.
- [9] J. Hofmeijer, M. C. Tjepkema-Cloostermans, M. J. van Putten, "Burst-Suppression With Identical Bursts: A Distinct EEG Pattern With Poor Outcome in Postanoxic Coma," *Clinical Neurophysiology*, vol. 125, no. 5, pp. 947-954, May, 2014.
- [10] B. T. Susam, et al., "Quantitative EEG Changes in Youth With ASD Following Brief Mindfulness Meditation Exercise," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 30, pp. 2395-2405, Aug, 2022.
- [11] P. J. Garcia-Laencina, G. Rodriguez-Bermudez, J. Roca-Dorda, "Exploring Dimensionality Reduction of EEG Features in Motor Imagery Task Classification," *Expert Systems with Applications*, vol. 41, no. 11, pp. 5285-5295, Sep, 2014.

Address for correspondence:

Nathan Riek
1238 Benedum Hall
3700 O'Hara Street, Pittsburgh, PA 15261
ntr14@pitt.edu