Using Embedding Extractor and Transformer Encoder for Predicting Neurological Recovery from Coma After Cardiac Arrest

Jan Pavlus, Kristyna Pijackova, Zuzana Koscova, Radovan Smisek, Ivo Viscor, Vojtech Travnicek, Petr Nejedly, Filip Plesinger

[1] Institute of Scientific Instruments of the CAS, Brno, Czech Republic

Abstract

This research presents a deep-learning framework designed to forecast neurological recovery following a cardiac arrest-induced coma. The framework is created by the team ISIBrno-AIMT as part of the Predicting Neurological Recovery from Coma After Cardiac Arrest: The George B. Moody PhysioNet Challenge 2023. Our approach involves a two-stage model: initially, the model derives low-dimensional embeddings from short electroencephalogram (EEG) segments (5 minutes), and subsequently, it combines the temporal progression (72 hours) of these embeddings to yield a comprehensive likelihood assessment of recovery outcomes. Regrettably, our submission was not evaluated in the ranking phase due to issues with the Docker pipeline.

1. Introduction

Cardiac arrest, characterized by the sudden loss of heart function, is a life-threatening event that often results in a profound lack of oxygen supply to the brain. This critical period of oxygen deprivation leads in 80 % of patients to a coma state [1]. Less than half of patients eventually wake up from a coma [2].

The prediction of neurological recovery from coma following cardiac arrest stands as a pivotal pursuit in contemporary clinical neurology. Many cardiac arrests occur outside the hospital and prognostic information from before the cardiac arrest is not available. The patient's prognosis is therefore estimated from data taken after the patient has fallen into a coma. The most widely used modality for assessing the severity of coma in clinical practice is electroencephalogram (EEG). Other important predictors of neurological outcome are based on clinical examination of motor response and ocular reflexes, and evaluation of certain biomarkers from blood serum, blood plasma, or CT imaging [3].

The disadvantages of continuous EEG analysis are subjectivity and time consumption. In addition, the evaluation can only be performed by neurologists with advanced training in neurophysiology. Therefore, fully automatic EEG classifiers have been created in recent years [4]. These classifiers have the potential to increase prediction accuracy while saving time for the human expert.

Utilization of deep-learning was proven to be effective in processing biological signals such as ECG [5] and EEG [6]. For this reason, we believe that deep-learning approach should be also effective for coma recovery prediction. This article aims to present a method for fully automatic prediction of neurological recovery from coma following cardiac arrest using deep learning models. This method was submitted to the George B. Moody PhysioNet Challenge 2023 [7].

2. Method

The method pipeline is depicted in Figure 1. Firstly, we preprocessed data by creating the bipolar signals from the uni-polar recordings in the dataset. All recordings were then resampled to 100 Hz and normalized by the z-score.

The model then comprises two primary components: the Embedding Extractor (Fig.1A) and the Temporal Model, both of which utilize extracted embeddings to analyze temporal patterns.

The Embedding Extractor (Fig.1A) is a convolutional neural network, primarily trained as a patient outcome classifier. It takes as input a spectrogram derived from a 5-minute segment of the electroencephalogram (EEG) signal. These 5-minute segments are extracted from the original 1-hour patient EEG signal recordings and are annotated with the corresponding patient outcomes. The network's output provides the probability of the patient's outcome. The signal embedding is obtained from the 128-dimensions linear layer of the embedding extractor.

The second component (Fig.1B) of the model is a neural network, which comprises a transformer encoder [8] followed by linear layers at the model's conclusion. The selection of the transformer encoder is motivated by its utilization of the attention mechanism, enabling it to

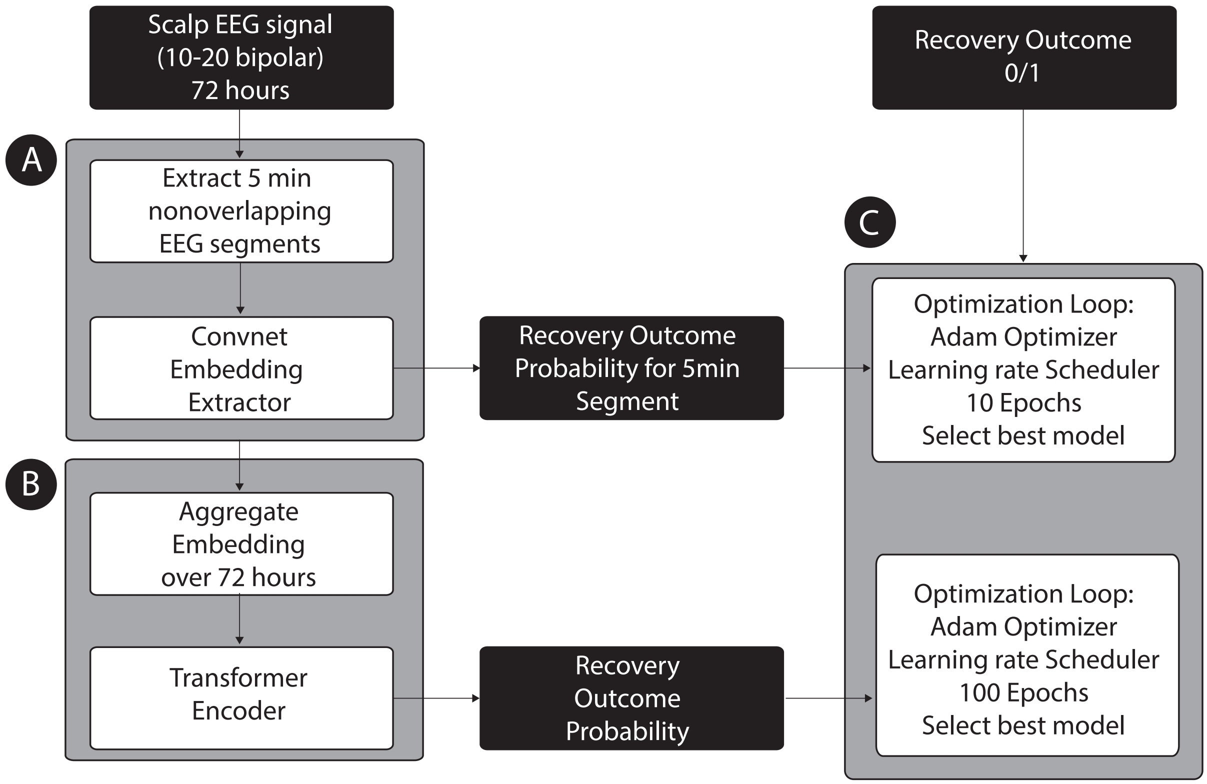

Figure 1. The figure shows the training pipeline and proposed model architecture i.e. convolutional neural network embedding extractor and transformer encoder for classifying recovery outcome. A) Embedding extractor from the 5-minute segments. B) Transformer encoder model that predicts outcome given also the temporal information. C) Optimization loops for A) and B) models. Each model is optimized separately.

consider the context from any part of the input. This model takes as input a matrix of embeddings that have been generated for each of the 72 hours of patient recordings as follows:

1. Temporal Segmentation: The data is divided into 5-minute chunks.
2. Embedding Extraction: These chunks are processed through an embedding extractor.
3. Embedding Aggregation: The mean and maximum embeddings are computed and concatenated.

These steps are applied to each non-missing recording within a 72-hour period for a patient. The resulting embeddings are organized into a matrix, with the x-axis representing time in hours. Missing hours are filled with vectors containing the value -1.

2.1. Transformer Encoder

The matrix is passed through a Transformer Encoder with three different linear layers, generating query (Q), key (K), and value (V) matrices. Multi-head attention is applied to Q, K, and V, where each head computes the attention mechanism using the formula:

$$Attention(Q, K, V) = softmax(\frac{QK^T}{\sqrt{d_k}})V \qquad (1)$$

where d_k represents the dimension of the key matrix K. The results from the multiple heads are concatenated and processed through another linear layer.

2.2. Training

The model is trained in two stages (Fig.1C). First, the embedding extractor is trained independently. Then, the network with the Transformer Encoder is trained, with the embedding extractor held fixed.

2.3. Prediction for Less Than 72 Hours

For predictions with less than 72 hours of data, the model treats the missing hours as if they were missing data and fills those hours with vectors containing the value -1.

2.4. ECG Signal Analysis:

In the investigation of ECG signals, we employed an analogous training process to that used for EEG data. Regrettably, the embedding extractor yielded inconclusive results when applied to the ECG signals. This ineffectiveness was primarily due to the absence of ECG recordings for a substantial number of patients and the suboptimal quality of the available data. Therefore, we did not include ECG data as input in the final submission.

2.5. Feature Examination

We also conducted an analysis of the provided features, such as patient Age and Gender, among others. Our examination revealed that the only noteworthy feature was the patient's age, which exhibited a discernible trend: older patients had a reduced likelihood of recovery (i.e., a poor outcome). During an unofficial experimentation phase, we explored training the model both with and without concatenating the features with the extracted embeddings. However, no discernible advantages were observed from including these features. Consequently, we excluded them from our model.

3. Results

Table 1. The table presented herein displays the outcomes of our cross-validation experiments, which were executed on training set. Specifically, the model was assessed by conducting tests on each individual hospital that had been omitted from the training set.

Hospital code	Percentage of public dataset	Percentage of poor outcomes	Challenge Score
A	43.00	53.26	0.410
B	19.77	72.00	0.151
D	13.67	67.47	0.518
E	12.19	79.73	0.780
F	11.37	60.87	0.619
Avg	–	–	0.500
Std	–	–	0.211

In light of our submission not attaining a ranking during the evaluation phase, we initiated a comprehensive experimental regimen, along with test-train splits. In this process, we conducted cross-validation individually on each of the hospitals contained within the dataset. Consequently, each iteration of the test set comprised solely one hospital, which had been deliberately excluded from the training set. Subsequently, we compiled and reported the results obtained for all participating hospitals, including the overall mean and standard deviation of the performance metrics.

It is important to note that only hospitals A, B, D, E, and F are accessible within the public section of the challenge dataset. Notably, hospital A accounts for half of the patients in the dataset [9].

All of our models underwent training from an initial state. The embedding extractor was trained utilizing a learning rate of 0.001, an exponential scheduler with a decay factor of 0.9, and the Adam optimizer. Concurrently, the temporal model was trained using a learning rate of 0.001, an exponential scheduler with a decay factor of 0.9, which was applied every 2 epochs, and the Adam optimizer. The outcomes of our experiments are presented in Table 1 and Table 2.

Table 2. True positive rate at a false positive rate of 0.05 (the official Challenge score) for our final selected entry (team ISIBrno-AIMT), including the ranking of our team on the hidden test set. We used 5-fold cross validation on the public training set, repeated scoring on the hidden validation set, and one-time scoring on the hidden test set.

Training	Validation	Test	Ranking
0.5 ± 0.2	-	-	-

4. Discussion

In the reported results in Table 1, hospitals D, E, and F displayed similar performance, as indicated by an average challenge score of 0.639 and a standard deviation of 0.108. These hospitals consistently produced the anticipated results.

However, hospital A exhibited a notably lower challenge score, standing at merely 0.410. It is worth noting that hospital A constitutes a substantial portion, approximately 43%, of the open dataset patient population. Consequently, the model was trained on slightly more than half of the dataset, which likely contributed to its reduced performance in this specific case.

The performance of the model on the hospital B dataset notably deviates from that observed in other datasets. This discrepancy could be attributed to the presence of specific features within the hospital B data that are absent in the datasets of other hospitals. Furthermore, it is conceivable that there may be inadequacies or shortcomings in the model training process that have contributed to this suboptimal performance.

5. Conclusion

This paper has introduced a deep-learning model for predicting neurological recovery in individuals who have experienced coma following cardiac arrest. We present the two-stage approach, involving the extraction of low-dimensional EEG embeddings from short signal segments (5 minutes) using a convolutional neural network and subsequent aggregation of temporal embeddings (over 72 hours) that are classified by transformer encoder network.

While our submission did not receive a ranking in the evaluation phase due to technical challenges with the Docker pipeline, our local environment experiments demonstrated the feasibility of the prediction task. We achieved a challenge score of 0.500 ± 0.211 (mean$\pm$std) through rigorous local cross-validation.

It's important to note that during our experiments, we found that the integration of ECG signals into our model did not lead to a notable improvement in the prediction scores. This unexpected observation underscores the complexity of the prediction task and the need for further investigation into the relationship between ECG signals and neurological recovery.

In summary, further refinement of our model and addressing the technical issues encountered during the submission/evaluation phase will be necessary. This research represents a step toward leveraging deep learning for improved neurological recovery prediction and underscores the importance of continued exploration in this critical healthcare domain.

Acknowledgments

The research was supported by the Czech Technological Agency grant number FW06010766 and the project RVO:68081731 by the Czech Academy of Sciences.

References

[1] Sandroni C, D'Arrigo S, Cacciola S, Hoedemaekers C, Westhall E, Kamps M, Taccone F, Poole D, Meijer F, Antonelli M, Hirsch K, Soar J, Nolan J, Cronberg T. Prediction of good neurological outcome in comatose survivors of cardiac arrest: a systematic review. Intensive Care Med 2022;48.

[2] Friberg H, Cronberg T, Dünser MW, Duranteau J, Horn J, Oddo M. Survey on current practices for neurological prognostication after cardiac arrest. Resuscitation 2015;90.

[3] Sandroni C, Cronberg T, Sekhon M. Brain injury after cardiac arrest: pathophysiology, treatment, and prognosis. Intensive Care Med 2021;47.

[4] Zheng WL, Amorim E, Jing J, Ge W, Hong S, Wu O, Ghassemi M, Lee JW, Sivaraju A, Pang T, Herman ST, Gaspard N, Ruijter BJ, Sun J, Tjepkema-Cloostermans MC, Hofmeijer J, van Putten MJ, Westover MB. Predicting neurological outcome in comatose patients after cardiac arrest with multiscale deep neural networks. Resuscitation 2021;169.

[5] Nejedly P, Ivora A, Smisek R, Viscor I, Koscova Z, Jurak P, Plesinger F. Classification of ecg using ensemble of residual cnns with attention mechanism. In 2021 Computing in Cardiology (CinC), volume 48. IEEE, 2021; 1–4.

[6] Nejedly P, Cimbalnik J, Klimes P, Plesinger F, Halamek J, Kremen V, Viscor I, Brinkmann BH, Pail M, Brazdil M, et al. Intracerebral eeg artifact identification using convolutional neural networks. Neuroinformatics 2019;17:225–234.

[7] Reyna MA, Amorim E, Sameni R, Weigle J, Elola A, Bahrami Rad A, Seyedi S, Kwon H, Zheng WL, Ghassemi M, et al. Predicting neurological recovery from coma after cardiac arrest: The george b. moody physionet challenge 2023. Computing in Cardiology 2023;50:1–4.

[8] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, Kaiser Ł, Polosukhin I. Attention is all you need. Advances in neural information processing systems 2017;30.

[9] Amorim E, Zheng WL, Ghassemi MM, Aghaeeaval M, Kandhare P, Karukonda V, Lee JW, Herman ST, Sivaraju A, Gaspard N, et al. The international cardiac arrest research consortium electroencephalography database. Critical Care Medicine 2023;10–1097.

Address for correspondence:

Jan Pavlus
Institute of Scientific Instruments of the CAS, v. v. i.
Kralovopolska 147
612 00 Brno
Czech Republic
xpavlu10@isibrno.cz