Detecting Harmful Brain Activity with Ensemble Learning

Ajay Pillay (apillay), Finn Moore (fdm), Nikhil Patel (nikhilap), Nilay Gautam (gnilay), Soham Das (sohamdas)

College of Engineering, University of Michigan

Introduction

This project seeks to address the challenge of accurately detecting and classifying forms of harmful brain activity using electroencephalography (EEG) signals from critically ill patients. Focusing on six distinct patterns, we aim to automate the analysis process, overcoming the limitations of manual review by specialized neurologists. The significance lies in advancing neurocritical care and epilepsy treatment through quicker and more accurate diagnoses. We propose a weighted ensemble model composed of three successful models: EfficientNet, ResNet, and WaveNet, significantly outperforming each model individually. Our model evaluation aligns with the competition's metric, utilizing Kullback-Leibler (KL) divergence. This project strives to contribute to accessible, accurate, and cost-effective neurocritical care solutions with transformative implications for medical treatments and diagnoses.

Data Preprocessing

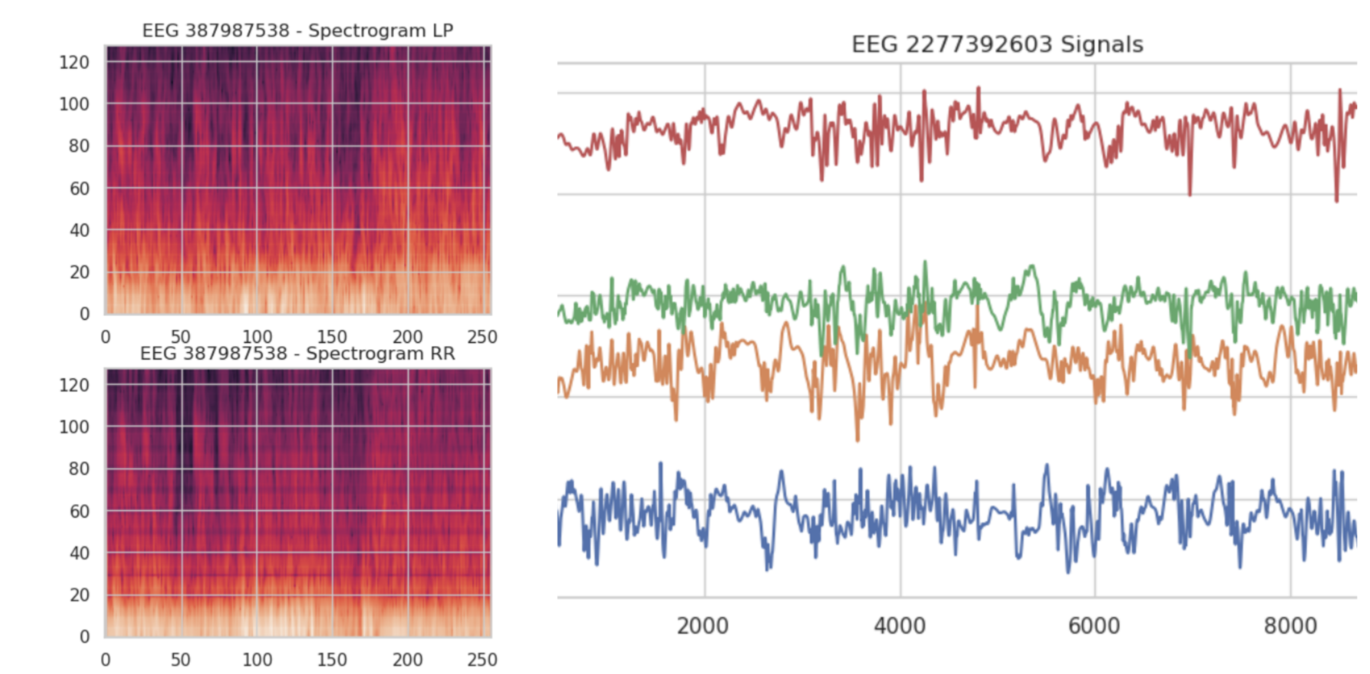

The EEG data used in this project was sourced from critically ill patients, with the aim of detecting harmful brain activity. The dataset comprised of raw EEG signals and their corresponding 11,138 spectrograms collected from 1,950 patients. Initial data handling involved cleaning and preprocessing the raw EEG data to ensure consistency and usability for machine learning applications.

- **Spectrogram Transformation:** Raw EEG signals were denoised and transformed into spectrograms using the Short-Time Fourier Transform (STFT). This conversion aids the application of CNNs by representing the data in the time-frequency domain.
- **Data Augmentation:** To increase the robustness of our model against overfitting and to enhance its ability to generalize across different EEG patterns, we applied data augmentation techniques such as horizontal flipping. This is particularly useful for image-based data like spectrograms as it simulated varying EEG conditions, thus expanding the diversity of training examples
- **Normalization and Standardization:** Both EEG signals and spectrograms were standardized by adjusting their scales to have means of 0 and a standard deviation of 1.
- **Filtering and Downsampling:** A Butterworth low-pass filter was applied to the raw EEG signals to attenuate high-frequency noise, enhancing the signal quality relevant for brain activity analysis. Raw EEG signals were also downsampled to reduce data dimensionality, making the computational processing more efficient without losing critical information.
- **Log Transformation and Scaling:** Log transformations were employed on spectrograms to improve the visibility of less prominent features by compressing the dynamic range of the signal intensities.

Method

In order to leverage the variety of provided data and features, we decided to employ an ensemble model comprised of three main models: EfficientNet, ResNet, and WaveNet. These are popular models that have proven to be successful at various other classification tasks. Our approach focused on tuning hyperparameters and strategically fusing these models together so that we could achieve a better performance than any single one of these models. The different models capture different underlying patterns in the feature space and combining the individual predictions would allow us to obtain an aggregated prediction, implicitly capturing the many underlying patterns, thus resulting in greater accuracy.

EfficientNet

EfficientNet is a CNN commonly used for image classification. It possesses a unique compound coefficient which is used to systematically scale up the network's depth, width, and resolution such that it achieves better accuracy and efficiency compared to scaling any singular dimension. We adapted the process described by Danial Zakaria [2] by employing transfer learning with EfficientNetB2 and additional non-frozen layers – a global average pooling 2D layer and a dense layer. Additionally, we trained EfficientNetB3 (containing more model parameters) on the original dataset.

ResNet

ResNet is a CNN that can have many layers – as it uses residual connections to mitigate the vanishing gradient problem. Thus, it is well suited for our time-series data, as we have 10,000 time steps for each sample. Within the network's depth, we include convolutional layers which learn local spatial patterns, and recurrent layers which model long-range temporal dependencies. We modified a version of ResNet (called EEGNet) developed by Med Ali Bouchhioua [1] by combining ResNet 1D Blocks in a sequential manner with different hyperparameters and passing the output through a fully connected layer to produce the final predictions.

WaveNet

WaveNet models long-range dependencies in waveform data very well. Using dilated convolutions, the kernel is spaced out, allowing it to have a larger coverage with the same number of parameters and computational cost, while preserving temporal resolution. When EEGs are administered, the electrodes are placed such that they form 4 montage chains that track the signals across different brain pathways. Grouping the data by chain and processing them in parallel proved usual, via a process adapted from Danial Zakaria [2].

Ensemble Pipeline

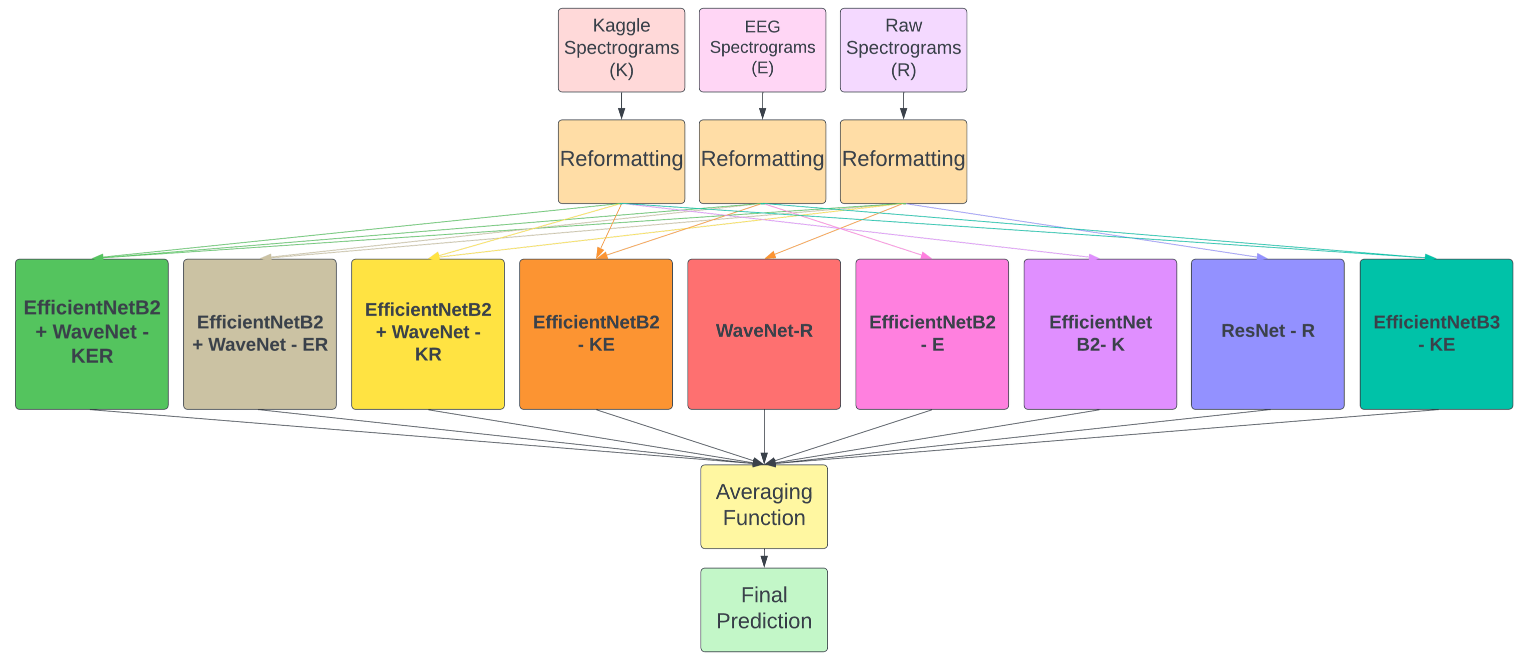

We had multiple datasets: Kaggle's spectrograms (K), EEG spectrograms (E), and raw EEG signals (R). Spectrograms are images while EEGs are time series waveform data. We built one ResNet model trained on EEG data, and experimented with various combinations of datasets and models as part of the ensemble, yielding 9 models in total. Inference was done by running the input through all 9 models and taking a weighted average of the resultant predictions.

Evaluation

KL divergence is a measure of how different two probability distributions are – in this case, it quantifies the difference between the predicted probability distribution and the observed target probability distribution determined by human experts. Formally, for discrete probability distributions P and Q on the same sample space $\mathcal{X}$, the KL divergence D_{KL} is defined as follows:

$$D_{KL}(P||Q) = \sum_{x \in \mathcal{X}} P(x) \log\left(\frac{P(x)}{Q(x)}\right)$$

This is undefined when $P(x) = 0$ and $Q(x) > 0$, so the convention adopted is to explicitly set the contribution to 0 in such cases. The goal is to minimize this score to model the real-world distribution as accurately as possible. We used a 5-fold cross-validation method during training for preliminary scoring.

Results

MODEL	DATA TYPE	LB SCORE Public	LB SCORE Private
ResNet	R	0.43	-
EfficientNetB2	K	0.41	-
EfficientNetB2	E	0.39	-
WaveNet	R	0.41	-
EfficientNetB2	KE	0.37	-
EfficientNetB2 + WaveNet	KR	0.39	-
EfficientNetB2 + WaveNet	ER	0.38	-
EfficientNetB2 + WaveNet	KER	0.36	-
Ensemble methods			
Initial Ensemble Inference		0.345	0.417
ResNet + Transformer		0.353	0.429
EfficientNetB3 Inference		0.350	0.421
Ensemble + ResNet + EfficientNetB3		0.335	0.408
KE + ER + KER + EfficientNetB3		0.321	0.390
KER + EfficientNetB3		0.318	0.386

Table 1. Models against LB Scores

Individual ensemble model scores on the leaderboard (LB) are detailed in Table 1. The final private LB score for our ensemble method was 0.389 ranking 912 (top 33%). Post-deadline, we got this score down to 0.386. Our results confirm the effectiveness of aggregating predictions across diverse models to enhance overall performance.

References

[1] Med Ali Bouchhioua.
HMS Resnet1D-GRU Inference — kaggle.com.
`https://www.kaggle.com/code/medali1992/hms-resnet1d-gru-inference`, 2024.
[Accessed 03-10-2024].

[2] Danial Zakaria.
Features+Head Starter — kaggle.com.
`https://www.kaggle.com/code/nartaa/features-head-starter/notebook`, 2024.
[Accessed 03-10-2024].