SMU Data Science Review

Volume 8
Number 2 *Summer 2024*

Article 4

2024

Applying Transfer Learning and Existing EEG Datasets to Identify Patients With ALS

Nibhrat Lohia
Southern Methodist University, nlohia@mail.smu.edu

Chris Mathew
Southern Methodist University, mathewc@mail.smu.edu

Garrett Shankel
Southern Methodist University, gshankel@mail.smu.edu

Recommended Citation
Lohia, Nibhrat; Mathew, Chris; and Shankel, Garrett (2024) "Applying Transfer Learning and Existing EEG Datasets to Identify Patients With ALS," *SMU Data Science Review*: Vol. 8: No. 2, Article 4.
Available at: https://scholar.smu.edu/datasciencereview/vol8/iss2/4

Applying Transfer Learning and Existing EEG Datasets to Identify Patients With ALS

Garrett Shankel[1], Christopher T. Mathew[1], Nibhrat Lohia[1]
[1] Master of Science in Data Science, Southern Methodist University,
Dallas, TX 75275 USA
garrett.shankel@mail.smu.edu
mathewc@mail.smu.edu
nlohia@mail.smu.edu

Abstract. Convolutional neural networks (CNNs) have helped medical practitioners re-design the diagnosis of electroencephalogram (EEG) recordings, resulting in a level of accuracy and efficiency that matches or exceeds experts in the field. This study aims to explore the application of VGG19, an algorithm based on CNN architecture, to diagnose small subsets of neurodegenerative diseases, specifically amyotrophic lateral sclerosis (ALS), a disease that reduces the nervous system's ability to control the muscles of the body. Since EEG recordings are scarce for ALS, this study will use a larger collection of EEG recordings taken on patients with more common neurological diseases, including but not limited to: Alzheimer's, seizures, and epilepsy. This learning was then transferred over to a smaller dataset containing recordings of ALS patients. The model was able to achieve an accuracy of 80% on the large collection and 68% on the ALS dataset. These results demonstrate that it is possible to use the transfer learning from VGG19 and achieve accurate results when applying it to smaller sets of EEG recordings related to different types of neurological diseases. The current methods used to diagnose ALS are either invasive or expensive. With the increased ability to diagnose using EEG recordings, early detection of ALS and better patient care can also be achieved.

1 Introduction

EEG recordings are scans that utilize electrodes on the scalp to measure electrical activity of the brain. They are used to diagnose brain related illnesses and conditions such as epilepsy, sleep disorders, and brain tumors. Currently, EEG readings are interpreted by hand and requires a skilled specialist such as a physician or neurologist to determine if an abnormality is present in the scan. This process can take time and is easily prone to human error. In many cases, it is possible that different specialists disagree in the diagnosis of a reading. Introducing an automated and accurate model to this process would reduce the burden on physicians and deliver faster, more reliable outputs for patients. This study will utilize scans taken over 15 years to develop a large-scale classification method that can be transferred to smaller datasets taken on ALS patients and will show that transfer learning can be a useful tool in diagnosing rare neurodegenerative diseases.

ALS poses many problems when it comes to diagnosis. Firstly, many of the methods are expensive or invasive, such as MRI's or spinal taps. The diagnosis can also be difficult due to the fact that many of the symptoms of ALS are similar to those of other neurological diseases. Diagnosis comes down to ruling out all other possibilities until ALS is the only disease left. Using EEG recordings is more cost effective when compared to MRI's. Also, using EEG recordings is easier for the patient than spinal taps. Along with improving patient care, a model allows researchers to perform feature analysis on EEG data and begin to pinpoint the signals that are unique to ALS.

One of the more popular approaches to classifying EEG recordings is utilizing neural networks, which are machine learning models inspired by the structure and function of the human brain. A more specific type of neural networks, CNNs, are used specifically for analyzing visual data, such as EEGs. CNNs are better equipped at handling spatial information compared to normal neural networks. This is done by employing convolutional filters that perform tasks such as edge detection for image classification problems. Methodologies attempting to automate EEG classification have resulted in a range of different accuracies. Most applications have been validated on single study groups and are able to achieve between 55% to 90% accuracy.

This research aims to use VGG19, a form of CNN, and other deep learning techniques to help automate the interpretation of ALS EEG readings. VGG19, an improvement on CNNs created by the Visual Geometry Group, makes use of 19 convolutional layers and a large body of images to "pretrain" the model. This training can then be leveraged in other image classification tasks. The capability of VGG19 to transfer its learning and its widespread use made it an obvious choice for this research.

In order to make use of EEG images in VGG19, the images must first be imported in their native form (*figure 1*).

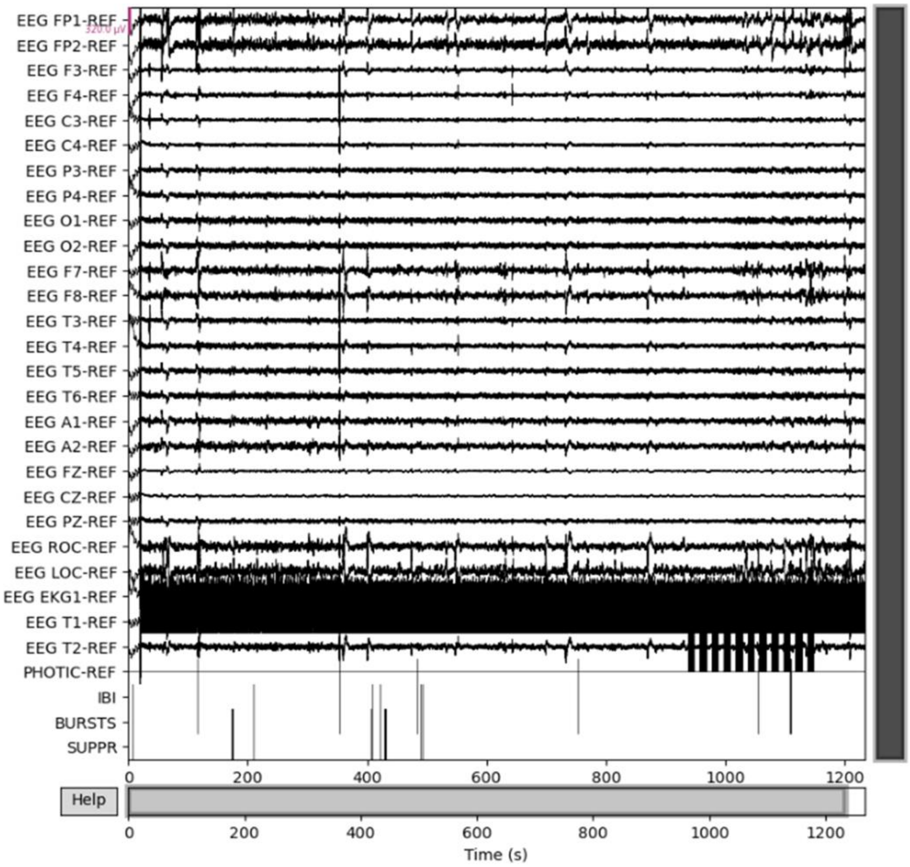

Fig. 1. An EEG reading produced using an .edf file from the TUH EEG Corpus. The readings represent voltage changes over time for each signal. The readings are broken into sections with lines one through four being the readings for the left sided chain. Lines five through eight being the right-side chain. Lines nine through twelve being the transverse chain. Lines thirteen and fourteen being the left and right eye electrodes, and the remaining lines being movements of the body.

Hidden within these separate readings are four different waveforms. Frequency will be used to differentiate between the different waveforms. First, there are delta wave forms with a frequency between 0-4 Hz. These are waveforms specific to slow wave sleep. If these are seen in a patient that is fully awake, then an issue is present. Second, there are theta waves between 4-8 Hz. These are typically related to drowsiness or slowed thinking while awake. Third, there are Alpha waves with a frequency between 8-13 Hz. These waveforms are associated with an individual that is awake with a normal functioning brain. Finally, there are Beta waves between 13-30 Hz which are related to stimulus or intense logical thinking. An EEG recording will be a mixture of these different readings.

This will not be the final form of the EEG recordings. The frequency readings must be converted into spectrograms. Spectrograms give a visual representation of the different frequencies present in each signal's waveform (*figure 2*). This allows for more information about the waveform to be collected in an image which can be easily fed into the VGG19 model.

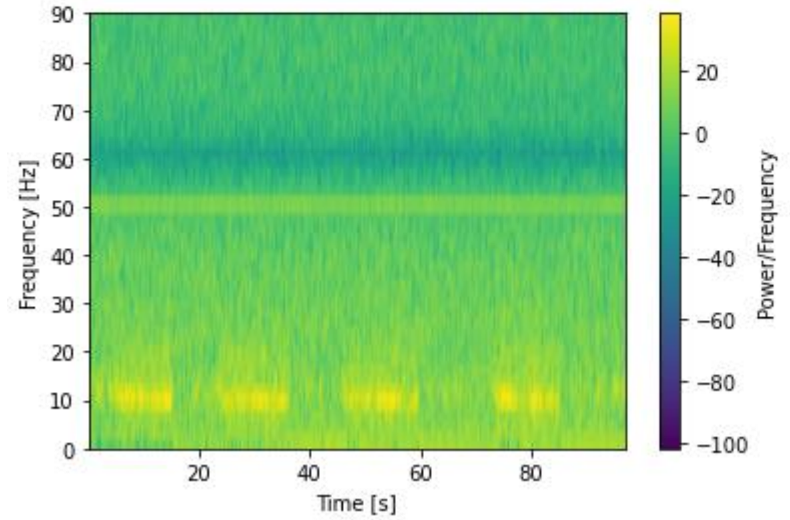

Fig. 2. A spectrogram showing component frequencies of an EEG signal. The colors represent the power of the frequency or the contribution of that single wave towards the overall waveform.

Several studies have shown that CNNs are able to accurately classify Alzheimer's, Parkinson's, seizures, and other neurodegenerative diseases. However, few studies focus on attempting to classify recordings of patients who have been diagnosed with ALS. ALS can be identified in patients who are already showing clear signs of the disease, but the degradation to the brain begins months to years before the diagnosis appears.

2 Literature Review

This literature review will be focused on four areas surrounding the problem of classifying EEG recordings: models that utilize EEG recordings, reducing noise in EEG signals for classification, current CNN research applied in seizure research, and ALS and the current status of collecting ALS EEG recordings.

2.1 Models Used for EEG Recording Classification

In a paper published by (Ferri, et al., 2020) the team poses multiple outcomes due to the research done on EEG and MRI recordings. The current paper will focus on and use the portion of Ferri's paper that deals with identifying Alzheimer's using machine learning. The team focused on developing the accuracy of Artificial Neural Networks by combining EEG and MRI recordings. The difference between accuracy when running on either set of data alone was increased significantly when running with both the MRI and EEG combined. This suggests that EEG alone may only be able to go so far in producing a diagnosis and other sources of data should be considered. If it is found that a CNN lacks accuracy when classifying neurodegenerative events with EEG alone, MRI scans may be incorporated.

A simpler approach would be a form of regression. Regression also serves as a reasonable baseline for the potential performance of more complex models. In the work by (Li et al., 2019) a graph regularized sparse linear regression model is applied to EEG recordings, with the goal being to identify the emotions of the subjects. They also used SVM and GraphSC as baselines to compare against the GRSLR model. The GRSLR model was able to achieve an accuracy of 67.97% which was between one to twenty percent more accurate than those used for comparison. While this does not directly correlate to neurological diseases, it shows that linear regression can be used accurately to read and interpret EEG recordings.

In the review published in (Craik et al., 2019) the Journal of Neural Engineering, it is shown that CNNs (or variants of CNNs) are frequently used when approaching EEG classification tasks. The CNNs and each variant also made use of a fully connected layer. CNNs were leveraged heavily across every major form of EEG classification. Emotion recognition, motor imagery, mental workload, seizure detection, sleep stage scoring, and ERP (Event Related Potential). None of the studies focused specifically on neurological diseases in EEGs.

Following close behind CNNs were Deep Belief Network (DBN) models. CNNs have been shown to outperform DBN's when tackling computer vision problems, making it surprising that a large amount of literature has made use of DBN's for EEG classification. It is possible that DBN's are favored in cases of noisy data or when EEG data is collected in abnormal scenarios. Since CNNs are highly dependent on their trained data they are weak to strong abnormalities. DBN's on the other hand are pre-trained on unlabeled examples. This could mean that the DBN models used for classification were able to learn patterns in the data that may have slipped past the technician that labeled the original data. If a highly trained human eye was fooled by an abnormal blip, it is highly likely that the CNN will be fooled as well. However, in practice CNNs outperform even on EEG tasks, but a combination of these two methods may give the best of both worlds.

In work done by Yuangfang Ren and Yan Wu (Ren et al., 2014) on unlabeled EEG data there is a promising hope for CDBN models. With this dataset the team was able to achieve accuracy scores around 85%, outperforming both CSP and MVAAR. It would have been good to benchmark this data against a CNN model, but without unlabeled data it may have been an apples and oranges comparison. The team's justification for using CDBN is compelling. CDBN was used here specifically to address the issues with classifying full images or high-dimensional data. Since high-dimensional data drives up computation, CDBN allows for the reduction of computations by sharing weights among the hidden and visible layers. CDBN, in the eyes of the researchers, is ideal for EEG data since it has several high-dimensional multi-channel data points.

It is important to note that most of the studies covered in this review used private data sources. This makes it difficult in practice to apply the accuracy values that they boast, without the ability to double check the data and its classifications.

2.2 Reducing noise in EEG signals for classification

One issue that solutions attempt to mitigate is noise that is contaminating the EEG signal. This noise can reduce the performance of classification models that are attempting to read the EEG recordings. As a result, noise reduction methods are employed to filter out unwanted noise while maintaining the integrity and information of the underlying signal.

One paper utilizes an approach called LAPPS (Least Absolute LP ($0<p<1$) Penalized Solution) (Bore et al., 2020). LAPPS is a specific type of granger analysis (GA), which tries to capture the causal relationship within a time series. This method is based on multiple variant auto-regression (MVAR). The paper used artificial noise generation to simulate more data for training and testing of the model. The steps were as follows: (1) Generate the 8-time series using the GCCA toolbox. (2) Contaminate the time series to model variations of noise. (3) Use different methods including LAPPS to estimate MVAR parameters and network parameters, followed by computing the MVAR bias. Compared to other methods of noise reduction tested in this study, LAPPS performed the best.

A study utilizing deep CNNs to classify epileptic states in EEGs notes in the paper that EEGs are notorious for having low signal-to-noise ratios. For this study, a wavelet threshold denoising method was utilized. The Daubechies wavelet of order 6 was chosen as the mother wavelet for discrete wavelet transformation or DWT (Gao et al., 2020.)

Another paper argues that noise contamination of the EEG reading is the number one reason for mediocre performance, which reiterates the importance of employing noise reduction techniques prior to training classification models (Klepl et al., 2022). In this paper, simpler filtering techniques were used to reduce noise. One technique used was a zero-phase 5^{th} order Butterworth filter. This was used to remove frequencies below 0.1 Hz and above 100 Hz. Another filter used was a zero-phase 4^{th} order Butterworth stop-band filter. This filter removed the 49-51 Hz range. The paper also noted that the data was sampled down to 250 Hz by employing an 8^{th} order Chebyshev type I filter. Similar simpler noise reduction techniques such as band-pass filters have been employed in other papers (Xu et al., 2019).

It is also important to note that certain parameters of the neural network can impact the model's sensitivity to noise. For example, LeakyReLU can be utilized over other activation functions to add robustness to noise (Acharya et al., 2018).

2.3 CNN In Seizure and Neurodegenerative Research

CNNs represent a type of neural network that makes use of convolution layers, pooling layers, and a fully connected layer all combined. The research done by (Abiyev et al., 2020) focuses specifically on the application of CNNs towards identifying seizure events in patients suspected to have lesions that need to be operated on. Their model utilized four convolutional layers, global average pooling, and a fully connected layer. This was all in combination with the RMSprop algorithm, specifically for addressing slow convergence. After training the CNN model, a comparative analysis was done using an SVM and NN model. The CNN model achieved a 96.67% accuracy rate, outperforming both models by ~21%.

Taking a step back to see how CNNs perform on a high level of brain activity classification, a team at the University of Missouri gave CNNs the task of classifying brain waves. As mentioned above, most CNNs are used to view an EEG as an image that will be classified as a specific event. These can be filled with noise making it difficult to interpret the results. Swapnil Joshi and their team set out to make sense of noisy readings using CNNs to interpret brain waves using LFP recordings (cortical tissue recordings of brain signals). These outperformed EEG recordings in the presence of pink noise and had better spatial and temporal resolution, reaching accuracy levels around 80% (Joshi et al., 2019). In addition, their team was able to identify unique brain waves with accuracies ranging from 76% to 90%.

The work performed by (Lawhern et al., 2018) applied CNNs in a slightly different approach. Brain computer interfaces or BCI's allow for the 'filtering' of EEG information into single categories of interface. Meaning certain events like vision-evoked potentials can be captured and identified by a BCI. In this work the team works to develop a CNN that can be applied to multiple different BCI's at once with limited training data. This is a large step forward in the EEG space. This shows that a CNN can be applied well even with limited training data. While there is no shortage of EEG data, the means for gathering it could soon be simplified with at-home monitoring.

Work performed on emotion recognition by (Wang et al., 2021) shows how a CNN, in their case FLDNet, can be used to filter out frames of an EEG based on a classification of ‘normal’ and ‘abnormal’. By classifying chunks of recordings as normal the performance and accuracy of the model can be increased. This would be a good first step in building a model by first classifying chunks of data. The chunks could be further refined to increase accuracy.

Most of the work performed when pairing deep learning and EEG recordings is focused either towards classifying events in a healthy brain (emotions, motor function, sleep states, etc.) or in identifying specific ailments like seizures or Alzheimer's. In work by Shu Oh and team, CNNs were applied to identifying Parkinson's disease in patients. Parkinson’s is of special concern to this paper since it is a disease that is marked with a slow degradation of the brain. In the Shu Oh study, twenty patients with PD and twenty healthy patients were observed. The CNN applied was able to reach an accuracy of 88.25%, a sensitivity of 84.71% and a specificity of 91.77% (Oh et al., 2020). The team accomplished this by using stratified tenfold cross-validation, Relu, softmax and Adam optimization. Each EEG recording lasted 5 minutes with the subjects moving as little as possible. These recordings were then broken into 1588 artifact-free epochs. The parameters for the CNN were tuned specifically for this group of people meaning it may be possible the model was overfitted on the data.

All these forms of research show how accurate CNNs can be with a large basket of mental illnesses, making it an ideal candidate to apply to the problem of identifying problematic EEGs of patients that may or may not have ALS.

2.4 ALS and the Collection of ALS EEG recordings

Amyotrophic Lateral Sclerosis (ALS) is a fatal neurodegenerative disease that causes muscular paralysis. The disorder affects both the upper and lower motor neurons, and it also begins to impair the ability to communicate when progressing to the advanced stages. While in the advanced stages, ALS is known as Locked-In Syndrome (LIS). In this state, the patient maintains cognitive function with little motor function. This includes conscious movement of the eyes.

One research paper documents the collection of EEGs and eye-tracking (ET) recordings from six ALS patients and 170 healthy patients (Le et al., 2024). The EEG recordings were obtained using a 32-electrode Emotive EPOC Flex device. The design of the experiment entailed nine different scenarios for each participant. Each scenario incorporated different tasks such as motor imagery, physical movements, and communication using the eye-tracking system. As for the methodology, rigorous calibration and quality control measures were employed to increase data accuracy. Preprocessing steps were taken by the users of the device to maintain the integrity of the raw data. The data was recorded over many sessions in the homes of the ALS patients. The data for the healthy patient was collected all at once in a laboratory setting. The goal of the paper was to present the dataset (EEGET-ALS) for use cases such as improving brain-computer interfaces (BCIs), study of motor cortex function, and the monitoring of motor impairment in ALS patients.

Preliminary experiments utilizing this dataset in tasks such as attention determination and person identification have reached accuracies of about 80% in certain cases. The dataset also contains detailed annotations and recording protocols, which make it a good resource for developing machine learning models such as convolutional neural networks.

3 Methods

3.1 Data

The data was sourced from the following resource:

https://isip.piconepress.com/projects/tuh_eeg/html/downloads.shtml. This dataset was modified to conform with HIPAA Privacy Rules by removing personal identifiers to make the data anonymous.

This corpus represents the largest known collection of EEG data collected and was brought together by the efforts of physicians and professors at Temple University Hospital and Temple University in Philadelphia, Pennsylvania. The data contains 26,846 clinical EEG recordings collected between 2002-2017.

The number of patients covered by the recordings is 14,987, with each patient averaging 1.79 sessions. 95% of the data is collected using the international 10-20 setup of electrodes on patients.

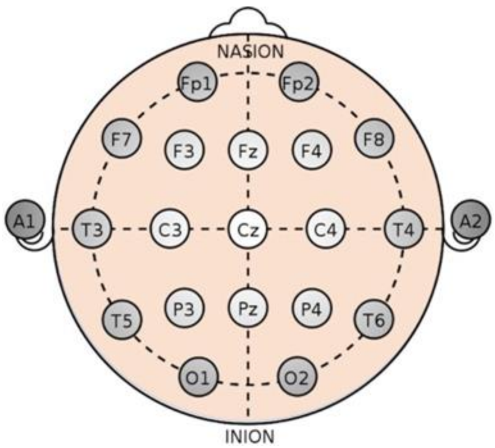

Fig. 3. The above diagram represents the international 10-20 setup. Each labelled circle represents a location where an electrode will be placed to measure brain activity. Even numbers denote the right hemisphere of the brain while odd numbers denote the left.

28 different independent variables are present in the dataset. 16 of these represent the 16 electrodes present in the international 10-20 setup. The remaining 12

of these are metadata that could or could not have been collected at the time of the recording including EKG recordings, photic stimuli, markers for specific types of infant brain activity, and extra electrodes placed along the scalp. Recordings are measured as voltage changes over time at each electrode. This results in a waveform as seen in *figure 4*.

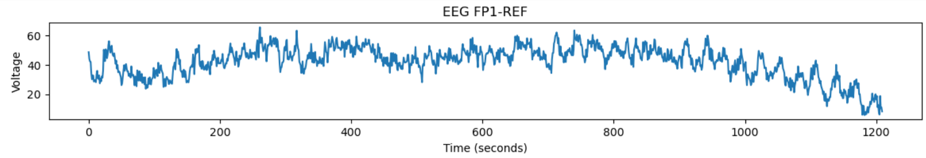

Fig. 4 A recording of a brainwave on the FP1 receptor. This is classified as normal brain activity. The recording represents voltage over time where voltage directly translates to the measured changes in voltage within the brain.

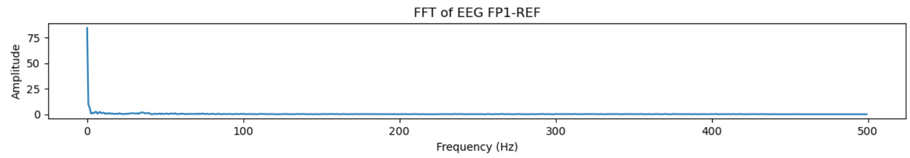

Fig. 5 Fast Fourier transform of a recording of a brainwave on the FP1 receptor. This is a translation of fig. 1 from the time domain into the frequency domain. Representing one of the constituent waveforms of the wave in fig. 1.

The signal variables in the data were run separately in a time domain as input to the model and as a frequency domain as input to the model. The time domain allowed for a benchmark of performance on the model. Converting the data into the frequency domain allowed for filtering and deeper analysis.

Fast Fourier transform (FFT) takes the waveform in the time domain and breaks it down into its constituent waveforms in the frequency domain. This means that all the building blocks of the wave in the time dimension are exposed. This enabled further analysis of what behaviors stand out among samples. Noise in a figure would suddenly gain distinct features that the model can be trained on.

The EEG recordings were converted from their native form (*figure 2*) into a spectrogram using FFT. FFT creates a spectrogram, a photographic representation of the constituent waveforms of the signals. This transforms the data into a form that can be fed into a CNN. From there the data was normalized using either a moving average or a form of local scaling. A lot of care needed to be taken to make sure seizure events were not “normalized” out of the data. Neurodegenerative events will naturally be “outliers”. Finally, the image data was converted into a grayscale image to further improve model performance (*figure 4*).

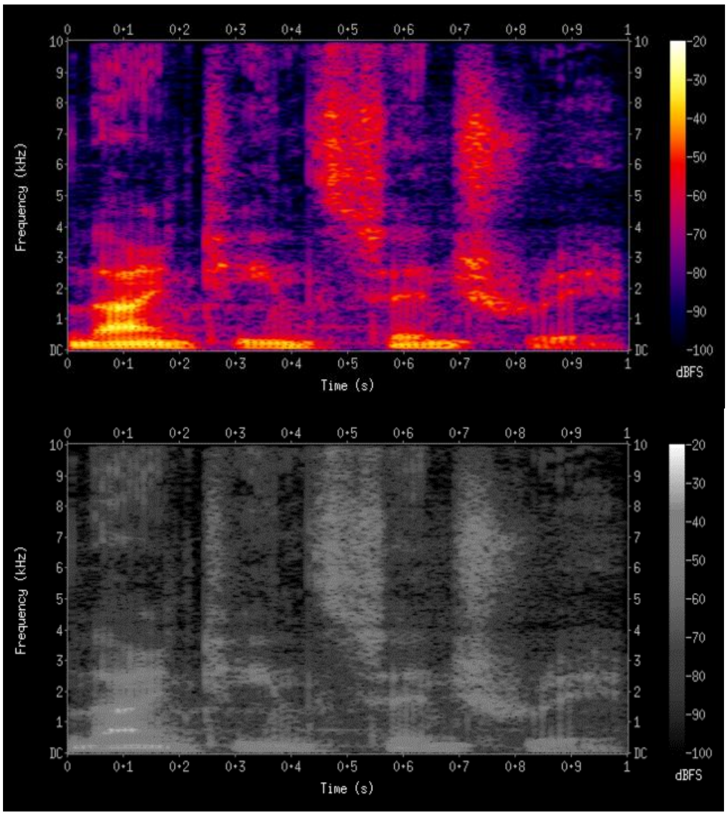

Fig. 6 Examples of a spectrogram and its greyscale conversion

Data from the study performed by Le and associates on ALS eye tracking was converted in the same way as the data from the TUH corpus. Once the model had been trained on the TUH corpus it was then transferred to the ALS dataset

3.2 Solutions Approach

This VGG19 model was trained on the TUH corpus. After tuning the model, the weights were frozen and transferred over to the ALS dataset. The major metrics that were focused on were accuracy and sensitivity. Specificity was not a primary focus of this model since false positives hold much lower risks when diagnosing a disease compared to the potential impact of a false negative. This accuracy was based on cross validation. Since this was a large dataset, only 5 folds were used when performing cross validation.

3.3 Implementation of the VGG19 Architecture

CNNs perform well on image data if the training datasets are large enough. In this case there was not a substantial number of images to train for both the TUH corpus and the ALS data. This issue led to the selection of the VGG19 algorithm so that previous image learning can be leveraged and transferred.

For this work the signals were averaged down from 27 to 3 so that data could be compressed and could better fit with the VGG19 architecture. From there, custom convolutional layers, dense layers, and dropout layers were added to further increase learning. Three convolutional layers were adjusted to have 32, 64, and 128 filters with each filter having a size of 3x3 and a 'relu' activation function. For each convolutional layer, batch normalizing was also used on the output of that layer to help speed up the training process. Next, the max pooling layers were adjusted to have a 2x2 filter. This helped simplify the input and aided the model in extracting the strongest features. Finally, two dense layers were applied: a layer with 128 neurons and a layer with 1 neuron. The final dense layer allowed the model to output either a 1 or 0.

Noise was a challenge with the EEG data. This caused the VGG19 model to have lower accuracy than expected. In order to overcome this, L1 and L2 regularization were implemented to help filter the data further. L1 and L2 regularization allowed for the generalization of the data, which helped improve accuracy when transferring the model to the ALS dataset. For the TUH corpus data there was an increase of roughly 3-5%, whereas for the ALS dataset, there was an increase of 1-2%.

4 Results

The performance of the VGG19 algorithm was evaluated using both training datasets: the EEG corpus from Temple University Hospital and the ALS dataset taken from the ALS eye tracking study. The results are summarized in *Table 1.* There was a noticeable decrease in the performance of the model on the ALS dataset, which was expected considering the small size of the dataset.

Table 1

Classification report comparing the results of the training on the general EEG data and the ALS data

	TUH General EEG			ALS EEG		
	Precision	*Recall*	*F1-Score*	*Precision*	*Recall*	*F1-Score*
Positive	82%	78%	75%	67%	62%	63%
Negative	78%	75%	77%	62%	59%	60%

Note: Positive represents those recordings of a normal patient. Negative represents those recordings of a patient with ALS

As seen above the precision indicates that our VGG19 model could achieve a rate of 62% positive prediction rate for the ALS dataset. Further iterations will seek to increase this value. As seen below in *Table 2,* the cross-validation results indicate that our VGG19 model achieved a mean accuracy of 80% with a standard deviation of 2.5%

on the general data and a mean accuracy of 68% with a standard deviation of 3.7%. Giving an accurate score to the application of VGG19 towards a small dataset of different neurological disorders.

Table 2

5-fold stratified cross-validation comparing the results of the training on the general EEG data and the ALS data

	TUH General EEG		ALS EEG	
	Accuracy	*Standard Deviation*	*Accuracy*	*Standard Deviation*
Percentage	80%	2.5%	68%	3.7%

5 Discussion

These results show that VGG19 can reasonably identify EEGs with ALS vs EEGs of normal recordings. This means that the transfer learning was a success in providing evidence that a model trained on more generic neurological disease data is a viable option when applied to more specific applications like ALS. Further research can expand on this by increasing public EEG data availability for ALS. From there, the accuracy can be improved to train a model with accuracies closer to a range of 80-90%. This can also be applied to feature extraction when the available ALS dataset has increased. The features identified using this model can be used to identify those waveforms that uniquely apply to ALS.

Caution should always be taken when using any model for diagnosis. Mislabeling a diagnosis can be dangerous for both the patient and the physician. Real world implementation of an algorithm such as this should not be taken lightly and should be met with strict testing and usage guidelines.

In the course of developing this model, it was interesting to see the improvements due to filtering and the final accuracy obtained. With such a small dataset, especially for a CNN, reaching reasonable levels of accuracy is difficult. Even the data itself was difficult to obtain. The concern for medical privacy is a valid one, but it makes public research challenging; however, the pursuit of research should not cost patients their anonymity. While the results are enlightening, personal data like this should always be handled with care to make sure the correct information is anonymized to protect people's privacy.

6 Conclusion

This study sought to explore the application of VGG19 towards the diagnosis of ALS. The study successfully demonstrated that transfer learning using VGG19 and a large corpus of EEG training data is a viable method in identifying EEG recordings related to ALS. The model achieved an accuracy of 80% on the larger dataset and 68% on the ALS dataset. This finding suggests that the VGG19 algorithm is effective in transferring its learning to ALS even in the presence of sparse data. This study also shows that VGG19 can noticeably improve the ability of clinicians to identify EEGs with evidence of neurological diseases.

The study was limited by the small amount of data available for ALS. CNNs require large amounts of data in order to effectively classify images. Even in the presence of sparse data, the model was still able to perform well but could see higher performance given more data to train on. For future research, it would be valuable to seek out more data on ALS and incorporate more data on abstract EEG recordings.

This study highlighted the potential of CNNs to help diagnose ALS in a non-invasive cost-effective way for patients. By improving diagnostic methods, and the features associated with certain diseases, healthcare can move closer towards a cure for difficult diseases.

Acknowledgments

We would like to express our gratitude to our advisor, Dr. Nibhrat Lohia, for his invaluable guidance, support, and encouragement throughout the development of this project. His insights and guidance have been instrumental in the successful completion of our project. We also extend our thanks to Temple University for providing access to the TUH EEG Corpus, which served as the foundation for our research. The availability of this extensive and well-documented dataset was crucial to the training and validation of our CNN model. This paper is a capstone project for the Southern Methodist University (SMU) Master of Science in Data Science (MSDS) program. We are grateful for the academic and technical resources provided by SMU, which have been essential to our work. Lastly, we appreciate the support from our peers and faculty members who have contributed to this project through constructive feedback and collaboration

References

1. Bore, J. C., Li, P., Harmah, D. J., Li, F., Yao, D., & Xu, P. (2020). Directed EEG neural network analysis by LAPPS (p≤1) Penalized sparse Granger approach. Neural Networks, 124, 213-222. https://doi.org/10.1016/j.neunet.2020.01.022

2. Lawhern, V. J., Solon, A. J., Waytowich, N. R., Gordon, S. M., Hung, C. P., & Lance, B. J. (2018). EEGNet: A compact convolutional neural network for EEG-based brain–computer interfaces. Journal of Neural Engineering, 15(5), 056013. https://doi.org/10.1088/1741-2552/aace8c

3. Ferri, R., Babiloni, C., Karami, V., Triggiani, A. I., Carducci, F., Noce, G., Lizio, R., Pascarelli, M. T., Soricelli, A., Amenta, F., Bozzao, A., Romano, A., Giubilei, F., Del Percio, C., Stocchi, F., Frisoni, G. B., Nobili, F., Patanè, L., & Arena, P. (Year). Stacked autoencoders as new models for an accurate Alzheimer's disease classification support using resting-state EEG and MRI measurements. Journal Name, Volume(Issue), Pages. https://doi.org/DOI

4. Wang, Z., Gu, T., Zhu, Y., Li, D., Yang, H., & Du, W. (2021). FLDNet: Frame-Level Distilling Neural Network for EEG Emotion Recognition. IEEE Journal of Biomedical and Health Informatics, 25(7), 2533-2544. https://doi.org/10.1109/JBHI.2021.3049119

5. Klepl, D., He, F., Wu, M., Blackburn, D. J., & Sarrigiannis, P. (Year). EEG-Based Graph Neural Network Classification of Alzheimer's Disease: An Empirical Evaluation of Functional Connectivity Methods. Journal Title, Volume(Issue), 2651-2660. https://doi.org/10.1109/TNSRE.2022.3204913

6. Gao, Y., Gao, B., Chen, Q., Liu, J., & Zhang, Y. (2020). Deep convolutional neural network-based epileptic electroencephalogram (EEG) Signal Classification. Frontiers in Neurology, 11. https://doi.org/10.3389/fneur.2020.00375

7. Acharya, U. R., Oh, S. L., Hagiwara, Y., Tan, J. H., & Adeli, H. (2018). Deep convolutional neural network for the automated detection and diagnosis of seizure using EEG signals. *Computers in Biology and Medicine*, *100*, 270–278. https://doi.org/10.1016/j.compbiomed.2017.09.017

8. Xu, G., Shen, X., Chen, S., Zong, Y., Zhang, C., Yue, H., Liu, M., Chen, F., & Che, W. (2019). A deep transfer convolutional neural network framework for EEG Signal Classification. *IEEE Access*, *7*, 112767–112776. https://doi.org/10.1109/access.2019.2930958

9. Abiyev, R., Arslan, M., John, B. I., Sekeroglu, B., & Ilhan, A. (2020). Identification of Epileptic EEG Signals Using Convolutional Neural Networks. *Applied Sciences, 10*(12), 4089. https://doi.org/10.3390/app10124089

10. Li, Y., Zheng, W., Cui, Z., Zong, Y., & Ge, S. (2019). EEG Emotion Recognition Based on Graph Regularized Sparse Linear Regression. *Neural Processing Letters, 49*(2), 555-571. https://doi.org/10.1007/s11063-018-9829-1

11. Joshi, S. R., Headley, D. B., Ho, K. C., Paré, D., & Nair, S. S. (2019). Classification of Brainwaves Using Convolutional Neural Network. *Proceedings of the ... European Signal Processing Conference (EUSIPCO). EUSIPCO (Conference)*, *2019*, 10.23919/eusipco.2019.8902952. https://doi.org/10.23919/eusipco.2019.8902952

12. Craik, A., He, Y., & Contreras-Vidal, J. L. (2019). Deep learning for electroencephalogram (EEG) classification tasks: A review. Journal of Neural Engineering, 16(3), 031001. https://doi.org/10.1088/1741-2552/ab0ab5

13. Y. Ren and Y. Wu, "Convolutional deep belief networks for feature extraction of EEG signal," 2014 International Joint Conference on Neural Networks (IJCNN), Beijing, China, 2014, pp. 2850-2853, doi: 10.1109/IJCNN.2014.6889383.

14. Oh, S.L., Hagiwara, Y., Raghavendra, U. et al. A deep learning approach for Parkinson's disease diagnosis from EEG signals. Neural Comput & Applic 32, 10927–10933 (2020). https://doi.org/10.1007/s00521-018-3689-5

15. Le, Thanh Ha; Ngo, Thi Duyen; Nguyen, Ba Hung; Kieu, Hai Dang; Nguyen, Minh Hoa; Nguyen, The Hoang Anh; et al. (2024). An EEG & Eye-Tracking Dataset of ALS Patients & Healthy People during Eye-Tracking-Based Spelling System Usage. figshare. Collection. https://doi.org/10.6084/m9.figshare.c.6910027.v1