

Interpretable Deep Learning Models for Single Trial Prediction of Balance Loss

Akshay Sujatha Ravindran^{1*}, Manuel Cestari¹ Christopher Malaya², Isaac John²,
Gerard E. Francisco³, Charles Layne² and José L Contreras Vidal¹

Abstract—Wearable robotic devices are being designed to assist the elderly population and other patients with locomotion disabilities. However, wearable robotics increases the risk from falling. Neuroimaging studies have provided evidence for the involvement of frontocentral and parietal cortices in postural control and this opens up the possibility of using decoders for early detection of balance loss by using electroencephalography (EEG). This study investigates the presence of commonly identified components of the perturbation evoked responses (PEP) when a person is in an exoskeleton. We also evaluated the feasibility of using single-trial EEG to predict the loss of balance using a convolution neural network. Overall, the model achieved a mean 5-fold cross-validation test accuracy of 75.2 % across six subjects with 50% as the chance level. We employed a gradient class activation map-based visualization technique for interpreting the decisions of the CNN and demonstrated that the network learns from PEP components present in these single trials. The high localization ability of Grad-CAM demonstrated here, opens up the possibilities for deploying CNN for ERP/PEP analysis while emphasizing on model interpretability.

I. INTRODUCTION

Falls are the leading cause of death, injury, and hospital admissions among the elderly population. Falls pose a significant threat not only to the safety and independence of seniors, but also lead to significant economic burdens. In 2018 alone, the direct costs from falls were estimated to be more than \$50B, with an average cost of \$30,000 per hospital visit [1][2]. There exists a need for fall prediction systems that can detect a loss of balance in seniors as early as possible, in order to lower the rate of morbidity-mortality.

Wearable robotic devices, such as powered exoskeletons, are being designed to assist the elderly population and other patients with locomotion disabilities (stroke survivors, spinal cord injured, cerebral Palsy etc.) [3]. Such devices are characterized by the implementation of traditional electric motors with the large gear reductions necessary to achieve high torques. However, these high torques come at the price of a reduced response velocity. Recovering a loss of balance involves quickly activating and engaging multiple muscle groups. Rapid joint activation is imperative to allow for an

adequate response time of the mechanical components of such wearable devices. Numerous applications of wearable devices to mitigate the risk of falling are being pursued for rehabilitative and assistance of elderly subjects and individuals with locomotion disabilities [4].

Neuroimaging (EEG) studies provided evidence for the involvement of frontal and parietal cortices in postural control [5][6]. It is possible that the neural information identifying a loss of balance can be detected prior to muscle activation with scalp electroencephalography (EEG) monitoring. The use of EEG scalp signals as markers for an imminent fall represents a novel approach to allow prompt balance compensation using wearable robotic devices. Early identification of fall-related EEG can be used to quickly activate the motors of the wearable robotic device in time to prevent the loss of balance. Human EEG studies regarding balance loss are rare and this lack of data is a major gap in our ability to understand how the brain detects a loss of balance and postural adjustments. These EEG signals provide crucial information regarding the activation time required for posture compensation or correction and may be used to actuate mechanical systems such as wearable robotics. Timely compensation in response to balance loss will prevent falls and their associated economic costs and health impact on elderly and locomotion impaired populations.

Prior studies have identified different components in the perturbation-evoked potentials (PEP) which are event-related potentials (ERP) evoked by different perturbations introduced either expectedly or unexpectedly [6][7]. These short latency responses are reported to be elicited by multisensory stimuli including visual, vestibular, and somatosensory. PEPs have several components consisting of an initial small positive wave (P1) followed by a large negative-going potential (N1). These are succeeded by a positive and negative wave (P2 and N2 respectively) [5].

In this study, initially we examine whether these components are preserved in healthy able-bodied subjects who wore exoskeletons during a series of postural perturbations. Traditionally, ERP components are identified after baseline correction and by averaging trials to improve the signal to noise ratio. However, for using EEG for predicting incoming falls in normal scenarios, there is a need to decode from single trial. Therefore, further, we evaluate the potential of

* Corresponding Author: Email - asujatharavindran@uh.edu

¹ NSF IUCRC BRAIN, Department of Electrical and Computer Engineering, Noninvasive Brain-Machine Interface System Laboratory, University of Houston, Houston, TX, United States of America

² NSF IUCRC BRAIN, Center for Neuromotor and Biomechanics Research, Department of Health and Human Performance, University of Houston, Houston, TX, United States of America

³ TIRR Memorial Herman Rehabilitation and Research Center, Houston, TX, United States of America

deep learning models to predict loss of balance from single trials by extracting information from these PEP components. Significant advancement has been made in using deep learning models as classifiers in EEG applications and has shown a significant advantage in improving the classification performance of the system. However, one of the major limitations of using deep learning models in EEG analysis is the lack of interpretability posed by these models [8][9][10]. However, to develop more confidence in using deep learning models, it is important to have model transparency, to better understand what aspects of the data the network is looking at to arrive at its decisions. This is a significant problem in sensitive applications such as detection of balance loss. Therefore, this paper puts emphasis on explaining the decisions made by the model evaluated using guided class activations maps (Grad-CAM) [11]. To the best of our knowledge, this is the first time Grad-CAM has been applied to deep learning applications on EEG to interpret model decisions.

II. METHODS

A. Participants

Six young healthy able-bodied subjects (4 males, 2 females) aged 18-32, participated in the study. All experimental protocols were approved by the Institutional Review Board (IRB) at the University of Houston. Written informed consent was obtained from each subject prior to the start of experimental procedures.

B. Experimental Setup

Subjects were instrumented with a 64-channel EEG cap, 8 Electromyography (EMGs) sensors, and body motion sensors (head, trunk, hip, thigh, knee, and shank). EEG data were recorded wirelessly at 250 Hz using active Ag/AgCl electrodes across the whole scalp (positioned according to a modified 10-20 international system) (BrainAmp DC, Brain Products, GmbH, Morrisville, NC). The electrodes corresponding to FT9 and FT10 based on the international 10/20 system were moved to replace AFz and FCz electrodes on the cap (corresponding to the ground and reference electrodes which were shifted to ear lobes). Similarly, channels TP9, TP10, PO9 and PO10, were used to record electrooculography (EOG) to measure eye-movement related artifacts during the experiments. Ten surface EMG electrodes from Delsys systems were mounted bilaterally on tibialis anterior, medial and lateral gastrocnemius, soleus, and peroneus longus. However, data from muscle sensors were not analyzed in this study.




Fig. 1. Experimental setup; Left: Subject standing on the neurocom system wearing the H2 exoskeleton; Right: Experimental protocol

Two basic conditions were tested, with or without a H2 robotic skeleton operating in the passive mode [12]. The order in which the exoskeleton was worn was randomized. In both conditions, the subject initially stood comfortably on the balance platform (Neurocom Balance Manager platform) for two minutes to attain baseline corticomuscular activity as well as the baseline ground reaction forces. The subject was then exposed to 32 backward translations of the support surface of constant duration (400 ms), displacement (6.35 cm) and velocity (15.875 cm/s). The timing of the perturbations was randomized so the subject would not be able to anticipate the time at which the perturbation would occur within the five second trial window. Each task was divided into two blocks of 16 perturbations separated by breaks to avoid fatigue. The protocol is summarized in Figure 1.

C. Signal Processing

All analyses were performed offline in MATLAB R2019b (MathWorks, MA) using the custom functions and the open-access EEG-processing toolbox, EEGLAB [13]. The pre-processing pipeline is summarized in Figure 2. The raw EEG and EOG signals from individual sets of trials were initially high pass filtered using a 4th order zero-phase Butterworth filter with a cutoff frequency of 0.1Hz. These were then passed into a low pass filter (4th order zero-phase Butterworth filter) with a cut off frequency of 50 Hz. The filtered signals from the individual sets of trials were then appended together. Eye artifacts were removed adaptively using an H-infinity filter [14] with γ parameter = 1.1 and q parameter set as $1e-11$. Artifact subspace reconstruction was then performed to remove large sudden bursts. A sliding window (length of 500 ms) and a less conservative threshold of 20-40 was used to identify corrupted subspaces. The signals were then decomposed into the independent components using Infomax independent component analysis (ICA). This was done to remove retained eye, muscle, or bundle artifacts. Stereotypical patterns of topoplot distribution, time-series data and PSD distribution, associated with the artifacts were evaluated manually to remove such artifactual components.




Fig. 2. Preprocessing flowchart

D. ERP Analysis

To identify the components in the ERP, initially, for each of the trials, baseline correction was performed by subtracting the mean of the one-second window prior to the perturbation onset for each individual channel. Then the data from the perturbation trials associated with the person being perturbed with and without the exoskeleton were averaged separately to increase the signal to noise ratio and view the ERP components.

E. Classification

A Convolutional neural network-based model was employed to predict the loss of balance from 200ms long single-trial windows.

Class 1: 200ms windows starting from 0.2 seconds prior to the perturbation onset until 500ms post perturbation with one sample (4ms) increment from a total of 64 trials (32 trials with and 32 trials without wearing the exoskeleton);

Class 0: To avoid potential bias instilled by comparing against baseline separated in time (potential drift/change in impedance could bias results), we selected windows based around these trials not involved with posture correction to serve as a control condition. Windows were selected one-second pre and post the perturbation onset and end. This also ensures there is higher EEG variability compared to continuous standing baseline conditions.

Initially, the trials were randomized and 70 percent of the trials were used to build the training set, 10% for validation, and 20% for the test set. This ensures, no single sample is common in either set and avoids any potential data leakages. This was repeated 5 times to create 5 fold cross-validation sets. In all the folds, both classes had an equal number of examples each.

F. Architecture




Fig. 3. CNN architecture block diagram

The architecture employed is shown in Figure 3. The input to the model was 200 ms long pre-processed EEG segment (50 samples x 60 channels). The data were normalized by dividing by the maximum absolute amplitude across the training set. The model consists of 4 convolution layers of 8 hidden units each (3x1 kernel dimension, stride=1). A final convolution layer with a kernel spanning all the channels (L5) was also used to combine information from all the channels. L2 kernel regularization was used with a value of 0.01 on all the convolution filters. Exponential linear unit was used as the activation function for all layers. Pooling layers (2x1 kernel dimension, stride = 2) succeeded two consecutive convolution blocks to reduce the dimension and improve the translation invariance associated with trial by trial variability.

This follows a Flatten layer prior to the two fully connected units with softmax activation function to generate prediction probability for each class. Dropout of 0.5 was used to reduce overfitting. Adam optimizer was used with a learning rate of 0.0001 to train the model with a batch size of 128 samples for a total of 250 epochs. Early stopping condition was used to stop training if validation loss stopped improving for 5 consecutive epochs. Different copies of the same model were used for each of the 6 subjects. The proposed models were implemented in python 3.6 using Keras 2.1.5 wrapper [15] with TensorFlow backend.

G. Grad-Class Activation Map

Grad CAM is a class discriminative localization technique used to identify discriminative regions in the input used by a CNN to arrive at the decision. The algorithm is detailed in Selvaraju et. al 2017 [11]. To obtain the localization map $L^c_{\text{Grad-CAM}} \in R^{u \times v}$ of height u and width v , we initially computed the gradient of the score for each of the class of interest with respect to the activation A_k of a layer (typically towards the end of the network). These are global averaged pooled over the height and width dimension to obtain importance weight α_k^c .

$$\alpha_k^c = \frac{1}{Z} \sum_i \sum_j \frac{\partial y^c}{\partial A_{ij}^k} \quad (1)$$

Here, Z is the number of samples/pixels in the feature map and i, j corresponds to the activation location. Then a weighted combination of the activation map, followed by a ReLU activation function was performed to obtain the activation map. ReLU was then employed to isolate features that have a positive influence on the class of interest while avoiding negatively contributing features.

$$L^c_{\text{Grad-CAM}} = \text{ReLU} \left(\sum_k \alpha_k^c A^k \right) \quad (2)$$

This resulted in a coarse heatmap of the same size as the convolutional feature map. Finally, we upsampled the heatmap to the input resolution. We used Keras-vis toolbox for implementing the grad-CAM [16]. In this study, the penultimate layer considered was P2 instead of L5 as we wanted to isolate relevant channels. After the model is trained, the mean activation map is then computed across the trials to evaluate the localization capability. From the averaged temporal CAM, the maximum value from each channel was estimated to identify the scalp distribution for the channels prioritized by the model for decoding.

III. RESULTS

A. ERP Analysis

Figure 4 shows the ERP components in the Cz channels from one of the representative subject (HS6).

As can be seen from the figure, P1, N1 P2 components, in particular, are present in both conditions with and without the exoskeleton. Here, the difference in amplitude was found to be indicative of ordering effect (EXO first or not) rather than resulting from wearing exoskeleton. However, the issue

TABLE I

5 FOLD CROSS-VALIDATION SCORES FROM ALL SUBJECTS; CHANCE LEVEL WAS 50% IN ALL CASES AND ALL NUMBERS ARE IN PERCENTAGES

Subject	Training Accuracy	Validation Accuracy	Test Accuracy	F score	P score	R score (%)
HS1	85.3 $\pm$ 2.6	78.1 $\pm$ 5.3	81.4 $\pm$ 3.7	81.4 $\pm$ 3.7	81.8 $\pm$ 3.8	81.4 $\pm$ 3.7
HS2	86.6 $\pm$ 1.8	80.7 $\pm$ 2.9	78.6 $\pm$ 5.5	78.5 $\pm$ 5.6	78.7 $\pm$ 5.5	78.6 $\pm$ 5.5
HS3	81.1 $\pm$ 3	76.8 $\pm$ 1.4	76.8 $\pm$ 5.1	76.8 $\pm$ 5.1	77.4 $\pm$ 5.7	76.9 $\pm$ 5.2
HS4	79.9 $\pm$ 0.8	75.8 $\pm$ 2.5	72.9 $\pm$ 5.3	72.8 $\pm$ 7	73.3 $\pm$ 5.5	72.9 $\pm$ 5.3
HS4	75.5 $\pm$ 4	66.8 $\pm$ 5.9	66.9 $\pm$ 6.9	66.8 $\pm$ 6.9	67.1 $\pm$ 7.2	66.9 $\pm$ 6.9
HS6	80.7 $\pm$ 2.5	72.1 $\pm$ 8.8	74.8 $\pm$ 4.4	74.7 $\pm$ 4.2	75.2 $\pm$ 4.3	74.8 $\pm$ 4.1
Total	81.5 $\pm$ 2.5	75.1 $\pm$ 4.5	75.2 $\pm$ 5.1	75.2 $\pm$ 5.1	75.6 $\pm$ 5.3	75.2 $\pm$ 5.1




Fig. 4. Averaged PEP from HS6. The arrows indicate the main PEP related components

of order effect is not addressed further in this paper. Since both conditions regardless of the order of EXO use contained EEG components of interest, the data from both conditions was combined for classification and model development.

B. Decoding Performance

The 5-fold cross-validation classification results are summarized in Table I. The mean across subject training classification score was 81.5 ± 2.5 %, validation score of 75.1 ± 4.5 and the test score was 75.2 ± 5.1 %. Subject 1 yielded the highest accuracy with a mean test accuracy of $81.4\% \pm 3.7\%$ whereas subject 5 had the lowest classification score of $66.9\% \pm 6.9$ % on the test set. The F1 score is similar to the test score indicating there exists no class imbalance issue and that the network is not prioritizing one class over the other

C. Interpreting the deep network

Here, the time-averaged Grad-CAM based relevance map is shown in Figure 5. It is able to temporally localize the relevant section corresponding to the N1 component and also isolate the most relevant channels. From the scalp distribution of the CAM, it is evident that the CNN looks at windows surrounding the N1 component, the CNN prioritizes the channels surrounding Cz to arrive at the decision. The model was able to identify relevant features in EEG in all the subjects both spatially and temporally.

To check whether the model only learns features from selected channels alone, ERP sections from different time points from a single subject (HS6) was sampled and how the Grad-CAM varies when the CNN sees different PEP components was evaluated. The result is summarized in Figure 5. Depending upon the data being shown, the CNN looks at different combinations of channels and time windows to arrive at the decision. Initially preceding the P1 component, close to the onset of perturbation, the Grad-CAM indicates that CNN looks at multisensory information from visual and

somatosensory regions of the brain. During windows with a clear P1 component, the relevance shifts towards central-parietal networks. During the N1 component, the CNN shifts more centrally towards Cz. Towards P2, the network looks mainly at the parietal regions. Overall, the attention/relevancy shifts depending on the prominent component present in the window.

IV. DISCUSSION

The usage of deep networks for sensitive applications requires a careful understanding of the features prioritized by the model to arrive at the decisions. Here we provided the first application of Grad-CAM based feature visualization of CNN on EEG and use it in the context of balance loss prediction. From the classification results, we validated that the model was able to predict the loss of balance using EEG data single trials of 200 ms and was able to localize (both spatially and temporally), the relevant ERP components which have been traditionally shown to be related to perturbation.

Among all the subjects, HS1 and HS2 have the highest classification score. They also have relatively localized relevancy maps both spatially and temporally. In contrast, HS5 had the lowest classification score. A careful exploration of the ERP component suggests the lack of a clear P1 peak for this subject. This subject had significant prior experience on the Neurocom platform and also had a relatively higher number of components associated with muscle artifacts compared to other subjects during ICA cleaning. Whether this might have influenced the results needs exploration.

The P1 component is said to be the earliest nonspecific cortical response to a perturbation that is driven by somatosensory input [5][6]. Comparing the localization corresponding to the P1 component, CNN looks both at central and parietal regions. This is in agreement with prior studies that identified the best localized cortical source that represents P1 to be Brodmann area (BA 5) which lies in the parietal region [17]. When the network sees windows having both P1 or N1 components, the network prioritizes both the sections temporally with higher relevance given to N1. Similar observations are seen during windows involving both N1 and P2 components as well. In both these situations when it localizes multiple regions temporally, the spatial spread is also higher as it involves central parietal on top of frontocentral regions. During windows of the P2 component,




Fig. 5. Grad-CAM results. Top row indicates the most relevant channels for the CNN to classify the window involving N1 component; Middle row indicates temporally which part of the input the network focussed on; Bottom row corresponds to the ERP of the input window with mean and shaded section corresponds to the standard deviation




Fig. 6. Grad-CAM results for different PEP components from HS6; Top row indicates the most relevant channels for the CNN to classify the respective input window; Middle row indicates temporally which part of the input the network focussed on; Bottom row shows the averaged ERP of the input window with mean and shaded section corresponds to the standard deviation

the CNN shifts towards the parietal and parieto-occipital regions.

The N1 component is one of the most reproducible components among the PEPs and has been previously reported to be maximally present in the FCz or Cz electrode in

prior studies [18][7]. In all subjects, the CNN model clearly identified Cz to be the most relevant channel to evaluate windows with the N1 component. The model was able to localize the component clearly temporally as well. Subject 3 had a larger spread in the scalp distribution; however

looking at the temporal relevancy map, we can see that the network also prioritized P1 peak in those windows, explaining the subject's larger spread towards the parietal regions. These inferences show that the network is looking at relevant features instead of inferring it from random data points. There is a possibility that the network is learning from movement/muscle artifacts, however after looking at the channels the network focused on, it appears that the network is prioritizing relevant brain regions instead of peripheral channels (which are expected to be easily corrupted by these artifacts).

In conclusion, in this research, we showed that the ERP components N1, P1, and P2 are present when subjects that wore exoskeleton were perturbed posturally. We also implemented a CNN architecture to decode balance loss from single-trial EEG. We introduced the application of Grad-CAM based visualization, to understand relevant features learned by CNN, highlighting its promise for future EEG applications. The Grad-CAM visualization technique showed that the model was able to identify and learn from different PEP components from single trials.

V. LIMITATIONS AND FUTURE CONSIDERATIONS

The emphasis of this paper was on the interpretability of CNN for ERP applications. It was not on finding the optimal model; therefore the classification performance could be improved with optimized hyperparameter selection or transfer learning. Since we observed the traditionally studied components to be retained even during EXO condition, we did not separate trials based on EXO vs No EXO to increase the number of samples for the model to learn from. However, there is a possibility of differences associated with these that could negatively affect the results. Currently, the decoding is performed using all 60 channels from 200 ms windows. However, looking at the grad CAM activity map, it clearly suggests the decoder looks at a narrower window and is very selective in the channels. This opens up the possibility for reducing the number of channels, time windows required, all of which are essential from a real-time application standpoint. Similarly, a key factor is the processing time and how early we can predict the incoming fall. Having an advance knowledge of an impending fall, provides the nervous system time to react to activate both simple reflexes and more complex automatic neuromuscular responses to prevent a fall. This might be true for some patient populations who have somewhat intact kinesthetic, and proprioceptive systems, that could produce intrinsic responses in combination with exoskeleton responses. However, individuals with serious spinal cord injuries are going to be reliant upon the exoskeleton responses, so kinesthetic and neuromuscular actions are going to have no meaning to them due to the nature of their condition. In such cases, it will be up to the brain to detect an impending fall with the hope that is enough lead time to activate the exoskeleton to prevent a fall. Future studies will explore these in more detail and should examine the changes involved during perturbed

walking as opposed to standing. This will likely involve more inter-cortical interactions and would be a useful test of this methodology's ability to identify those interactions.

ACKNOWLEDGMENT

This project was sponsored by the NSF IUCRC Building Reliable Advances and Innovation in Neurotechnology (BRAIN) Center Award (NSF Award 1650536)

REFERENCES

- [1] R. Rajagopalan, I. Litvan, and T.-P. Jung, "Fall prediction and prevention systems: recent trends, challenges, and future research directions," *Sensors*, vol. 17, no. 11, p. 2509, 2017.
- [2] C. S. Florence, G. Bergen, A. Atherly, E. Burns, J. Stevens, and C. Drake, "Medical costs of fatal and nonfatal falls in older adults," *Journal of the American Geriatrics Society*, vol. 66, no. 4, pp. 693–698, 2018.
- [3] S. Viteckova, P. Kutilek, and M. Jirina, "Wearable lower limb robotics: A review," *Biocybernetics and biomedical engineering*, vol. 33, no. 2, pp. 96–105, 2013.
- [4] Y. He, D. Eguren, T. P. Luu, and J. L. Contreras-Vidal, "Risk management and regulations for lower limb medical exoskeletons: a review," *Medical devices (Auckland, NZ)*, vol. 10, p. 89, 2017.
- [5] J. P. Varghese, R. E. McIlroy, and M. Barnett-Cowan, "Perturbation-evoked potentials: significance and application in balance control research," *Neuroscience & Biobehavioral Reviews*, vol. 83, pp. 267–280, 2017.
- [6] R. A. Ozdemir, J. L. Contreras-Vidal, and W. H. Paloski, "Cortical control of upright stance in elderly," *Mechanisms of ageing and development*, vol. 169, pp. 19–31, 2018.
- [7] A. L. Adkin, S. Quant, B. E. Maki, and W. E. McIlroy, "Cortical responses associated with predictable and unpredictable compensatory balance reactions," *Experimental brain research*, vol. 172, no. 1, p. 85, 2006.
- [8] I. Sturm, S. Lapuschkin, W. Samek, and K.-R. Müller, "Interpretable deep neural networks for single-trial eeg classification," *Journal of neuroscience methods*, vol. 274, pp. 141–145, 2016.
- [9] R. T. Schirrmester, J. T. Springenberg, L. D. J. Fiederer, M. Glasstetter, K. Eggensperger, M. Tangemann, F. Hutter, W. Burgard, and T. Ball, "Deep learning with convolutional neural networks for eeg decoding and visualization," *Human brain mapping*, vol. 38, no. 11, pp. 5391–5420, 2017.
- [10] A. S. Ravindran, A. Mobiny, J. G. Cruz-Garza, A. Paek, A. Kopteva, and J. L. C. Vidal, "Assaying neural activity of children during video game play in public spaces: a deep learning approach," *Journal of neural engineering*, vol. 16, no. 3, p. 036028, 2019.
- [11] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-cam: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 618–626.
- [12] M. Bortole, A. Venkatakrishnan, F. Zhu, J. C. Moreno, G. E. Francisco, J. L. Pons, and J. L. Contreras-Vidal, "The h2 robotic exoskeleton for gait rehabilitation after stroke: early findings from a clinical study," *Journal of neuroengineering and rehabilitation*, vol. 12, no. 1, p. 54, 2015.
- [13] A. Delorme and S. Makeig, "Eeglab: an open source toolbox for analysis of single-trial eeg dynamics including independent component analysis," *Journal of neuroscience methods*, vol. 134, no. 1, pp. 9–21, 2004.
- [14] A. Kilicarslan, R. G. Grossman, and J. L. Contreras-Vidal, "A robust adaptive denoising framework for real-time artifact removal in scalp eeg measurements," *Journal of neural engineering*, vol. 13, no. 2, p. 026013, 2016.
- [15] F. Chollet *et al.*, "Keras," <https://keras.io>, 2015.
- [16] R. Kotikalapudi and contributors, "keras-vis," <https://github.com/raghakot/keras-vis>, 2017.
- [17] A. Mierau, T. Hülsdünker, and H. K. Strüder, "Changes in cortical activity associated with adaptive behavior during repeated balance perturbation of unpredictable timing," *Frontiers in Behavioral Neuroscience*, vol. 9, p. 272, 2015.
- [18] S. Quant, A. L. Adkin, W. Staines, and W. McIlroy, "Cortical activation following a balance disturbance," *Experimental brain research*, vol. 155, no. 3, pp. 393–400, 2004.