

Exploring Relationships between Cerebral and Peripheral Biosignals with Neural Networks

Alexander H. Hatteland*
ETH Zürich
Email: hattelaa@student.ethz.ch

Ričards Marcinkevičs*
ETH Zürich
* shared first authorship

Renaud Marquis
CHUV Lausanne

Thomas Frick
IBM Research Europe

Ilona Hubbard
CHUV Lausanne

Julia E. Vogt
ETH Zürich

Thomas Brunschwiler†
IBM Research Europe

Philippe Ryvlin†
CHUV Lausanne
† shared last authorship

Abstract—Autonomic peripheral activity is partly governed by brain autonomic centers. However, there is still a lot of uncertainties regarding the precise link between peripheral and central autonomic biosignals. Clarifying these links could have a profound impact on the interpretability, and thus usefulness, of peripheral autonomic biosignals captured with wearable devices.

In this study, we take advantage of a unique dataset consisting of intracranial stereo-electroencephalography (SEEG) and peripheral biosignals acquired simultaneously for several days from four subjects undergoing epilepsy monitoring. Compared to previous work, we apply a deep neural network to explore high-dimensional nonlinear correlations between the cerebral brainwaves and variations in heart rate and electrodermal activity (EDA). Further, neural network explainability methods were applied to identify most relevant brainwave frequencies, brain regions and temporal information to predict a specific biosignal.

Strongest brain-peripheral correlations were observed from contacts located in the central autonomic network, in particular in the alpha, theta and 52 to 58 Hz frequency band. Furthermore, a temporal delay of 12 to 14 s between SEEG and EDA signal was observed. Finally, we believe that this pilot study demonstrates a promising approach to mapping brain-peripheral relationships in a data-driven manner by leveraging the expressiveness of deep neural networks.

I. INTRODUCTION

With the latest advancements in sensing devices and machine learning, it is possible to gather and analyse increasing amounts of information about human activity at various levels [1]. Exploring connections between brain and peripheral physiological signals could improve the utility of peripheral biosignals at health monitoring [2] and capturing the state of the human brain.

© 2021 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Published version: Hatteland, A. H., Marcinkevičs, R., Marquis, R., Frick, T., Hubbard, I., Vogt, J. E., Brunschwiler, T., & Ryvlin, P. (2021). Exploring Relationships between Cerebral and Peripheral Biosignals with Neural Networks. In *2021 IEEE International Conference on Digital Health (ICDH)* (pp. 103-113). IEEE. doi: 10.1109/ICDH52753.2021.00022

Previous studies focusing on the relationships between brain and peripheral biosignals have either utilised conventional statistical models, such as multiple linear regression [3], or have outlined their own mechanistic models supported by experimental findings [4], [5].

In this paper, we present a machine learning pipeline for analysing the relationship between stereo-electroencephalography (SEEG) and peripheral signals, *i.e.* heart rate and electrodermal activity (EDA). It relies on convolutional neural networks (CNNs) [6] to find dependencies between high-dimensional time-series data and to infer complex relationships between the brain and peripheral signals acquired from the SEEG electrodes and an Empatica E4 wearable wristband [7]. To improve interpretability of our analysis and elucidate inferred relationships, we utilise gradient-based and remove-and-retrain network explanation techniques [8]. Our approach allows identifying brain regions most closely related to different bodily functions as well as disentangling the contributions of specific features within these signals. To the best of our knowledge, this pilot study is a first attempt to utilise deep neural networks to model brain-peripheral relationships.

In Section II of this paper, we provide the background on time-series prediction, explainability methods, neural networks for EEG data, and previous studies focusing on brain-peripheral relationships. The experimental setup with the resulting dataset and the analysis pipeline are explained in Section III. The findings of this study are provided and discussed in Sections IV and V.

II. BACKGROUND & RELATED WORK

A. Time-series Prediction Using Neural Networks

In our analysis, we investigate the relationship between peripheral and SEEG signals by training and validating predictive neural-network-based models. Currently, deep neural networks are ubiquitous in machine learning applications [9], including time-series prediction and forecasting [10]. Conventional neural network approaches to modelling and predicting time-series data include recurrent neural networks (RNN) [11], long short-term memory (LSTM) [12], and temporal convolutional networks [13]. They allow capturing patterns in high-dimensional

time-series with complex, nonlinear dynamics. For instance, Ahmed *et al.* [14] compared various machine learning time-series forecasting techniques empirically, including regression trees, support vector regression, Gaussian processes, and multilayer perceptrons. They found that neural networks alongside with Gaussian processes were able to discover more complex relationships and hence, performed better at forecasting.

B. Neural Network Explainability Methods

While deep neural networks excel at finding complex relationships between input features and the output [15], their significant drawback when compared to classical time-series models [16] is that they are “*black boxes*” [17], [18]. In particular, it is difficult to disentangle which features a neural network uses to predict a specific output value [17]. For our analysis, it is essential to identify the most relevant features and to understand how the input and output signals are linked. Herein, we provide a brief overview of relevant techniques that are utilised to “*explain*” predictions made by neural network models.

1) *Perturbation-based Methods*: Perturbation-based explanations provide an attribution for input features by masking or altering inputs and quantifying the resulting change in the output [19], [20], [21]. Such approaches are commonly referred to as “*occlusion*”. They allow measuring the importance of different input features in a neural network *post hoc* without retraining or adjusting the model.

2) *Gradient-based Methods*: Gradient-based methods [8] utilise neural network gradients to examine the influence of hidden and input units by backpropagating through the network. They assign an importance score to each neuron for different outcomes. For example, DeepLIFT [17] is a backpropagation-based method that uses a neutral baseline input, *e.g.* an all-black image, to compare neuron activations against the reference value resulting from the given baseline.

3) *Remove-and-retrain*: Another approach to quantifying the importance of input features is to compare networks trained on different subsets of the original features [22]. By retraining the network and comparing the performance of different models it is possible to identify features that are crucial for predicting the output [22].

C. Neural Networks for EEG data

While no prior work has explored a general relationship between SEEG and peripheral signals captured from wearables using deep neural networks, applications of neural networks to EEG data are abundant. Most classification tasks for EEG signals fall into six categories [23]: sleep stage scoring [24]; seizure detection [25], [26], [27], [28]; emotion recognition [29], [30], [31]; mental workload estimation [32], [33]; and other applications [34], [35], [36]. The latter include depression screening [36], Alzheimer’s disease detection [34], and gender classification [35].

Prior work utilises various EEG signal features and representations [23]. Some approaches focus on the raw signal [37] or averages of several signals [25]. It is also common

to extract handmade features [38] or to combine them with the raw signal. Another approach is to convert the EEG time-series to 2D image-like representations, *e.g.* spectrograms [24], colour scale topology [39], Fourier feature maps [40] or 2D/3D correlation grids [41]. Spectrograms are arguably the most prevalent choice among 2D representations of the EEG signal [23]. While representations of input signals differ across various architectures and tasks, the frequency range that the models focus on tends to be between 0.5-60 Hz [23].

Model architectures vary across different feature extraction strategies, including deep belief networks [42], recurrent neural networks [43], stacked autoencoders [44], multilayer perceptrons [45], CNNs [24], [46], and hybrid architectures [47], which combine multiple approaches.

Particularly relevant to our work is the Sleep Phase Identification with Neural networks for Domain-invariant LEarning (SPINDLE) pipeline, proposed by Miladinović *et al.* [24]. SPINDLE uses 2D convolutional filters in a CNN to extract relevant features from preprocessed EEG spectrograms. In this paper, we combine SPINDLE’s preprocessing and its architecture with the AlexNet network [48] to extract relevant features from SEEG data.

D. Brain-peripheral Relationships

There exists a rich body of literature on the relationships between brain and peripheral biosignals. Herein, we provide a summary of biological background and findings relevant to our analysis.

1) *Central Autonomic Network*: The central autonomic network (CAN) is a brain region in charge of the behavioural responses essential for survival, as shown experimentally by Benarroch [4]. CAN consists of the insular cortex, amygdala, hypothalamus, periaqueductal gray, parabrachial complex, the nucleus of the solitary tract (NTS), and the ventrolateral portions of the medulla [4].

2) *Heart Rate*: In particular, the insular cortex and amygdala were found to have the strongest autonomic control of the body, including the cardiac responses [4]. This was further confirmed by Thayer *et al.* [5]. Their findings show that there are two regions of the brain that have a relationship with the heart rate’s emotional responses: the cingulate cortex and the ventromedial prefrontal cortex. Furthermore, there is increasing evidence for a connection between the insular cortex and cardiac responses [49], [50]. To sum up, previous studies show that changes in heart rate are linked with activity in the CAN area of the brain. However, an exact form of this relationship and relevant cerebral signal features remain still unknown.

3) *Sweat Response*: Fredrikson *et al.* [3] investigated the forebrain relationship with sweat level responses. They found that electrodermal activity was related to the cingulate cortex, secondary visual cortex, and right inferior parietal cortex, with the strongest relationship to the right insular cortex, which is a part of the CAN region. Furthermore, Farrell *et al.* [51] found that there are several areas linked to thermogenic sweating, including the insula, anterior and posterior cingulate cortex,




Fig. 1. Schematic of the experimental setup including all recording devices worn by patients. The CT scan depicts SEEG electrode placement for patient 1. The peripheral biosignal recordings were made by the Empatica E4 wristband [54] placed on patient’s non-dominant hand.

premotor cortex, thalamus, lentiform nuclei, and cerebellum. Other areas were more prevalent during psychogenic sweating events: midcingulate cortex, parietal cortex, premotor cortex, occipital cortex, and cerebellum. Other studies have also suggested links between changes in electrical conductance of the skin and the insula [52], [53].

III. METHOD

In this section, we provide a brief summary of the dataset and outline the complete data analysis pipeline applied to the SEEG and peripheral data.

A. Experimental Setup & Data Overview

The data includes intracranial SEEG (DIXI Medical, France) recorded at Lausanne University Hospital (CHUV), Switzerland. Four epileptic patients underwent epilepsy monitoring during several days prior to brain surgery, and SEEG electrodes were placed in the suspected seizure onset zones specific to each patient. On top of the cerebral monitoring, a wearable wrist device (E4, Empatica, US) was worn to capture peripheral biosignals. A schematic of the overall experimental setup is shown in Figure 1. Table I provides an overview of the characteristics of the four patients enrolled in the study. Age, gender, and BMI were not controlled for when choosing subjects. All patients gave written informed consent, and the research was approved by the local ethical committee (ref. 2020-01818). The recordings lasted between three to seven days and include both day- and nighttime. During the recordings, patients were allowed to move freely in the hospital. Since we focus on the *general* relationship between the

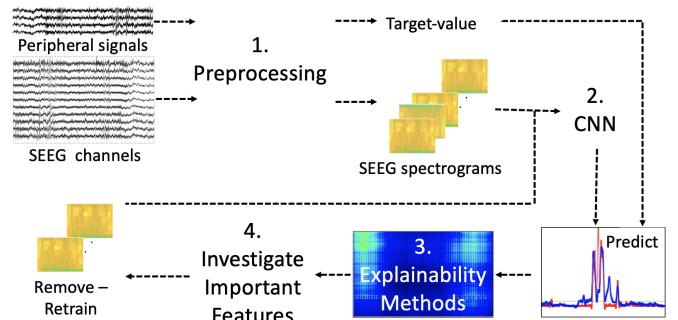


Fig. 2. The analysis pipeline for exploring relationships between SEEG and peripheral signals in patients undergoing intracranial EEG recordings.

peripheral and cerebral biosignals, we only consider regular inter-ictal periods.

Patient	Age, yr	Gender	BMI, kg/m ²
1	41	Male	39.2
2	18	Female	20.4
3	40	Male	23.5
4	26	Female	24.7

TABLE I
OVERVIEW OF THE PATIENTS ENROLLED IN THE STUDY.

Each of the SEEG electrodes has several contact points (see Figure 1). These contact points measure and record the electrical potential in the brain. They are numbered from 1 to N , where 1 corresponds to the deepest contact in the brain and N is closest to the skull, with the prefix indicating the specific electrode. One electrode usually penetrates multiple brain-regions, therefore, it is important to account for the location of the specific contact points, rather than performing the analysis in an electrode-wise manner.

The peripheral signal measurements came from an E4 wristband (Empatica [54], US) located on the patient’s non-dominant hand. It was replaced by a new fully charged wristband every 24 hours. Patients were instructed to wear the wristband at all times, except when they had to take it off, for instance, during shower.

B. The Analysis Pipeline

The analysis pipeline consists of four stages shown in Figure 2: (i) the preprocessing of the respective signals; (ii) the neural network training and validation; (iii) application of the explainability methods to identify important features; (iv) and finally, attributing identified relationships to relevant brain regions. Our holistic approach can unravel nonlinear interactions between the signals that classical time-series analysis methods, *e.g.* see [16], would not detect and at the same time, allows exploring these relationships further by explaining neural network predictions.

We consider two peripheral biosignals as the target variables: (i) heart rate and (ii) electrodermal activity. We treat both of them independently. They serve as the target variables which the neural network predicts using a spectrogram-based

representation of the SEEG data. Thus, we perform supervised learning for a time-series regression task.

1) *Preprocessing*: Both the SEEG recordings and the peripheral signals contain unwanted biological artefacts, which originate from the patient internally, as well as the ambient environment. Therefore, preprocessing is crucial for the downstream analysis. All of the preprocessing described below is performed in the MATLAB (v.2020a) computing environment.

a) *Artifact Removal & Preprocessing of the SEEG Signals*: In our analysis, we focus on five frequency bands of EEG signals: *delta* (0.5-4 Hz), *theta* (4-8 Hz), *alpha* (8-13 Hz), *beta* (13-30 Hz), and *gamma* (30-100 Hz) [55].

We down-sample the SEEG signal from 1024 to 256 Hz by applying a low-pass filter to prevent aliasing (fifth-order Butterworth filter with a cut-off frequency of 128 Hz). This is followed by a cascaded integral comb with compensation filter [56] and decimation (down-sampling). Henceforth, the SEEG sampling rate will be 256 Hz, unless mentioned otherwise. A notch filter is applied at 50 Hz to remove the line noise.

Since this analysis focuses on the general relationship between the cerebral and peripheral biosignals, time instances where the behaviour is abnormal, *e.g.* ictal phases as labeled by physicians, are discarded. To ensure that the data does not contain any outliers due to the pre- and post-ictal phases, segments 5 minutes prior and 20 minutes after each annotated ictal phase were removed as well.

Following the approach of SPINDLE [24], the SEEG time-series is converted into a time-frequency representation by applying the discrete short-time Fourier transform (STFT) [57]. For each channel, the STFT is computed using a two second (512 samples) periodic Hamming window with a step size of 125 milliseconds (32 samples) on the down-sampled SEEG signal. A Hamming window is used to reduce the edge effects [24]. In the time-frequency domain, the squared magnitude of the STFT is taken yielding the power spectral density (PSD). The PSD is further log-transformed and then standardised for each frequency to zero-mean and unit-variance.

Multiple standardised frequency windows are concatenated together producing a feature map that describes frequencies over several seconds. These signals are split into 16-second intervals, containing the PSD for every frequency from 0.5 to 64 Hz (with a frequency resolution of 0.5 Hz). These feature maps or images are stacked (one feature map per contact point) before they are fed into the CNN. Although the resulting images contain the same information as the 2D signal, we have found that it is easier for the CNN to extract relevant features from the time-frequency representations.

b) *Artifact Removal & Preprocessing for Peripheral Signals*: During data acquisition, the patients were instructed to keep the wristband on them at all times. However, there were exceptions to this. For example, the Empatica E4 is not designed to withstand water [54], so patients were allowed to remove the wristband when necessary. When the wristband was off, the temperature recording from the thermophile sensor dropped to within 2°C of ambient temperature (22°C), while the acceleration remained constant. We remove all of these

time periods using the information from the temperature and acceleration sensors.

One of the peripheral signals we consider is the electrodermal activity. The EDA signal refers to the variation of the skin's electrical conductance (SC), which changes depending on the patient's sweat level. SC time-series can be characterized by a slowly varying tonic and a fast-varying phasic activity. To investigate the relationship between the brain and the instantaneous skin conductance changes, the raw EDA signal is preprocessed to remove the slow varying tonic part [58]. The tonic proportion is estimated using nonnegative deconvolution [58] followed by smoothing with a Gaussian window. After that the phasic signal is found by subtracting the tonic proportion from the original SC signal. These steps are performed using the skin conductance analysis software developed by Ledalab (v. 3.4.9) [59].

Finally, the preprocessed measurements of the peripheral signals at the end of each SEEG time-window are used as the target variable for the CNN (see Figure 3). This way, the network is given SEEG information from up to 16 seconds prior to the peripheral signal it tries to predict.

2) *Convolutional Neural Network*: By converting the cerebral biosignals to time-frequency representations it is possible to leverage 2D CNN architectures. We propose a CNN architecture, inspired by SPINDLE [24] and AlexNet [48]. By training and validating this neural network we investigate whether it is possible to predict peripheral signals using SEEG. While there exist more performant neural network architectures for time series regression, *e.g.* RNNs and LSTMs, in this analysis, we focus on assessing relationships between signals, rather than achieving the best possible predictive performance. Moreover, it is essential that the chosen architecture is amenable to the application of the explanation techniques mentioned before.

The overall architecture of the proposed network has the same number of layers as AlexNet (8 layers). The layers are modified to better fit features from spectrograms and the different input dimensionality. The model is adjusted to take N signals as inputs, depending on the number of specific contact points implanted in a patient. The first layers of the network are taken from the SPINDLE's CNN architecture [24].

Each of the N contact points is represented as one input spectrogram. Thus, the input dimensions are $128 \times 128 \times N$. The first two dimensions represent time (0-16 seconds prior to the predicted signal value) and frequency (0.5-64 Hz), respectively. The network consists of four convolutional layers, followed by four fully connected layers that predict a normalised value between 0 and 1. A schematic of the network is shown in Figure 3. In images, neighbouring pixels often belong to the same object. This is not the case in spectrograms, where the different frequencies are often non-locally distributed [60]. Therefore, similar to previous networks operating on spectrograms [24], [61], [62], our CNN has smaller kernels in the first layers, since more extensive areas of feature retrieval are unnecessary.

Unlike most convolutional networks, our model has a rela-




Fig. 3. The network architecture, based of AlexNet [48] and SPINDLE [24]. The overall architecture is adapted from AlexNet, while the first two layers are similar to SPINDLE. Green layers represent max-pooling, while the red represent 2D convolutions. The blue layers are fully connected. Time is introduced as an additional feature to each of the fully connected layers, to account for the circadian rhythm.

tively large input space, with N channels, N ranging from 69 to 217 across patients. The significantly larger input size makes it unfeasible to use only one convolutional layer, as proposed in SPINDLE [24]. A large amount of input spectrograms also makes it necessary to increase the number of filters compared to the traditional AlexNet architecture. The filter sizes are then decreased in the last convolutional layer to make the feature space in the fully connected layers smaller. Furthermore, time of the day is introduced as an additional feature into each of the fully connected layers, to account for SEEG patterns specific to the circadian rhythm.

The convolutional layers (shown in red in Figure 3) implement m filter operations of size 3×3 . After the filter, the rectified linear unit (ReLU) activation is applied to all neurons independently. Max-pooling layers down-sample the features by only choosing the maximal value of the non-overlapping rectangles of size 2×2 . Batch normalisation and dropout layers are added to prevent overfitting. The time features, introduced in the fully connected layers, are encoded as $\cos(\cdot)$ and $\sin(\cdot)$ of the day time, to result in a differentiable periodic signal. Finally, mean squared error (MSE) loss function is used when training the network.

The model is implemented in Python (v.3.8.3) programming language using the PyTorch (v. 1.8) machine learning library [63].

3) *Training & Testing*: To test whether an association between the input SEEG and the output peripheral signal is significant, we evaluate the networks predictive performance on unseen data. The dataset is split temporally into training and test sets, 80% and 20%, respectively. Unlike in standard train-test split, training and test sets are kept as separate temporal intervals to avoid overly optimistic results due to a strong correlation between training and test data points, for a detailed discussion see [64]. We train and validate a separate model for each subject and adjust the architecture to account for the different input sizes due to the differing numbers of contact points across subjects. To better assess the performance of our models across *multiple* time periods [65], we conduct K -fold cross-validation with $K = 5$ separate train-test splits.

The neural networks are trained for 10 epochs using Adam

optimiser with a mini-batch size of 16 and a learning rate of 0.0001. We did not observe further improvements in performance on withheld data from longer training. These parameters are tuned manually on withheld data. The parameters are tuned to achieve the highest R^2 (the coefficient of determination) score [66]. The coefficient of determination usually ranges between 0 to 1, where 1 indicates a perfect model fit. We use this metric to assess the models fit and the strength of the relationships between inputs and outputs.

4) *Identifying Important Features & Brain Regions*: To identify which features and brain regions are most closely related to peripheral biosignals we use several explainability methods. In the time-frequency domain, the DeepLIFT [17] method is applied to produce attribution maps for all frequencies and time intervals of the spectrogram. As mentioned before, DeepLIFT requires providing a baseline sample. During our experiments, different baselines were tested yielding similar attribution maps, therefore, a value of 0 is used as a baseline for all spectrogram features. Due to the architecture of our neural network, it is not possible to directly use DeepLIFT attributions to quantify the importance of entire brain regions. Therefore, we resort to the remove-and-retrain approach by removing subsets of contact points corresponding to relevant regions of the brain, such as hippocampus, white matter, and insula.

IV. RESULTS

As mentioned before, our analysis focuses on training neural network models to predict the heart rate and EDA based on SEEGs and validating these models to establish whether learnt dependencies are significant. In the following subsections, we provide the results for the two peripheral biosignals.

A. Relationship between SEEG and Heart Rate

Table II contains R^2 scores achieved by the model in 5-fold cross-validation, for both the raw and low-pass filtered heart rate signal. The low-pass peripheral signal is extracted using a fifth order Butterworth filter with a cut-off of 25 seconds. As can be seen, the model predicts the heart rate of patients 1 and 4 with a much higher R^2 score than for patients 2 and 3. Note, that the patients with contact points in the CAN

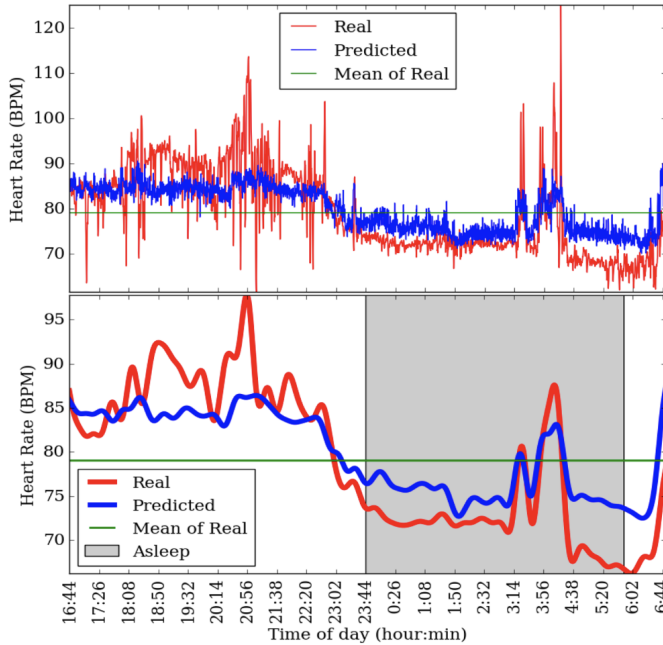


Fig. 4. Predicted vs. actual heart rate for patient 1. The *top* graph shows the raw signals, while the *bottom* graph contains low-pass filtered signals.

region consistently have an R^2 score greater-than 0. In fact, the score is significantly greater-than 0 for patients 1, 3, and 4 (t -test $p < 0.025$). For patient 2, we observe no significant relationship between SEEG and the heart rate (t -test $p > 0.05$). In general, the model can predict slow changes in heart rate much better than short excursions, as signified by the larger low-pass score.

Patient	CAN	R^2 score (\pm SD)		
		Full signal	p	Low-pass
1	✓	0.268(\pm 0.187)	0.016	0.436(\pm 0.236)
2	✗	0.052(\pm 0.141)	0.228	0.120(\pm 0.228)
3	✓	0.093(\pm 0.047)	0.006	0.298(\pm 0.139)
4	✓	0.285(\pm 0.197)	0.016	0.454(\pm 0.279)

TABLE II

CROSS-VALIDATED R^2 SCORE FOR THE FOUR PATIENTS WHEN PREDICTING THE HEART RATE BASED ON SEEG SIGNALS. THE MODEL IS TRAINED AND TESTED ON FIVE DIFFERENT TIME INTERVALS (80% TRAIN, 20% TEST). R^2 SCORE IS MEASURED ON THE RAW DATA (“full signal”) AND ON THE LOW-PASS FILTERED TARGET AND PREDICTED SIGNALS. STANDARD DEVIATIONS (SD) ARE TAKEN ACROSS 5 TEST FOLDS.

Figure 4 shows the predicted against ground truth peripheral signal. For both the raw and low-pass filtered signals, predictions appear to agree with the ground truth values. Similar observations can be made for the other patients. Since our model does not have any memory of earlier time stamps, overall, we conclude that it is possible to predict patient’s approximate heart rate purely based on the instantaneous SEEG signal. This implies that the SEEG signals are associated with instantaneous heart rate changes.

When inspecting the R^2 score across different test folds (see standard deviations in Table II), it can be seen that the




Fig. 5. Residuals vs. fitted values for the heart rate prediction in patient 1. Colours indicate the wakefulness status of the patient. Green denotes data points labelled *awake* and blue denotes data points labelled *asleep*.

model performs considerably worse on some folds (hence, a large SD). We observe that the model performs worse when the subjects are awake. It is possible that subject’s wakefulness is a confounder in our analysis and some of the performance of our models might be attributed to detecting the wakefulness status. Nevertheless, our models still manage to retrieve relevant information for the patients with contact points in the CAN region and produce somewhat reliable prediction during the daytime period alone.

The latter finding is aligned with the residual analysis (see Figure 5). Tests sets with data from both sleep and wakefulness periods show that the predictions form two distinct clusters correlated with the sleep/wakefulness state of the patient. Those states are derived using the peripheral device’s accelerometer by detecting the periods when the patient remained still. This is cross-referenced with the other sensors. Note, that our wakefulness/sleep labels are merely a proxy.

Using explanation methods described before, we can explore which features are utilised by the network to predict the heart rate. Figure 6 shows most influential input features when predicting the heart rate from the SEEG spectrograms. The figure was produced for patient 1, but similar attribution maps were observed for all other patients. The most relevant features reside in the alpha and beta frequency ranges, while some useful information can also be found in higher frequencies (around 55 Hz).

Brain region explanations are derived by intrasubject model performance analysis through retraining of the network (see Figure 7). The results further confirm that the network performs best with signals originating from the CAN region.

In particular, the network performs well when only using the information from the insula (anterior, middle, and posterior gyrus), despite the low number of contacts in that brain region, as indicated in Figure 7.

B. Relationship Between SEEG and Electrodermal Activity

We now explore the relationship between SEEG and the phasic activity in EDA. The low-pass peripheral signal is again




Fig. 6. Focus of the network when predicting the heart rate from the spectrogram for patient 1. The attribution map is produced using DeepLIFT [17] and averaged across the whole testing period. A baseline of 0 is used for all inputs.




Fig. 7. R^2 score obtained by the network when retraining on specific contact points located in different brain regions.

extracted using a fifth order Butterworth filter with a cut-off of 25 seconds. Table III provides an overview across all patients for predicting the phasic driver of the EDA signal. Interestingly, a similar tendency is reflected in these results as for the heart rate (*cf.* Table II). Patients 1 and 4 have a stronger association between the SEEG and EDA with relatively higher R^2 scores. Indeed, for these patients the cross-validated R^2 score is significantly greater-than 0 (t -test $p < 0.025$).

Patient	CAN?	R^2 score (\pm SD)		
		Full signal	p	Low-pass
1	✓	0.425(\pm 0.106)	0.001	0.397(\pm 0.111)
2	✗	-0.369(\pm 0.700)	0.152	-1.891(\pm 3.784)
3	✓	0.001(\pm 0.262)	0.497	-0.104(\pm 0.740)
4	✓	0.205(\pm 0.156)	0.021	0.338(\pm 0.219)

TABLE III

CROSS-VALIDATED R^2 SCORE FOR THE FOUR PATIENTS WHEN PREDICTING THE PHASIC DRIVER OF EDA BASED ON SEEG SIGNALS. THE MODEL IS TRAINED AND TESTED ON FIVE DIFFERENT TIME INTERVALS (80% TRAIN, 20% TEST). R^2 SCORE IS MEASURED ON THE RAW DATA (“full signal”) AND ON THE LOW-PASS FILTERED TARGET AND PREDICTED SIGNALS. STANDARD DEVIATIONS (SD) ARE TAKEN ACROSS 5 TEST FOLDS.




Fig. 8. Predicted vs. actual phasic driver signal for patient 1. The *top* graph shows the raw signals, while the *bottom* graph contains low-pass filtered signals.

Figures 8 and 9 shows the full and low-pass filtered predictions vs. ground truth signal of the phasic driver (EDA) for patients 1 and 2. For patient 2 (average $R^2 < 0$), the model is unsuccessful at predicting large EDA peaks which correspond to sweat glands opening. For patient 1 (average $R^2 = 0.425$), however, network’s predictions follow the magnitude of the phasic driver comparatively well. In this case, the model is clearly able to predict when the the sweat glands open and the overall strength of the phasic driver.

We again examine which SEEG features are most dominant in predicting the EDA signal. Clearest patterns could be observed for patient 1, for whom the model exhibits good performance. Most relevant features lie in the alpha and beta frequency ranges, similarly to the most dominant features in predicting the heart rate (Figure 10). Additional relevant features can be observed in the higher frequency band around 55 Hz. Furthermore, feature relevance is time dependent, indicating a time delay between the SEEG and EDA signals, where the phasic driver of the EDA signal lags from the SEEG by around 12-14 seconds. This is very different from the time-invariant attribution map obtained for the heart rate (*cf.* Figure 6).

By retraining the model on specific contact points in patient 1, we identify strongest performance achieved by contact points located in the insula (CAN) (see Figure 11), similarly to the heart rate response. The highest test R^2 score is achieved when using the middle and posterior gyrus of the insula. In this patient, the hippocampus is the most epileptic region which

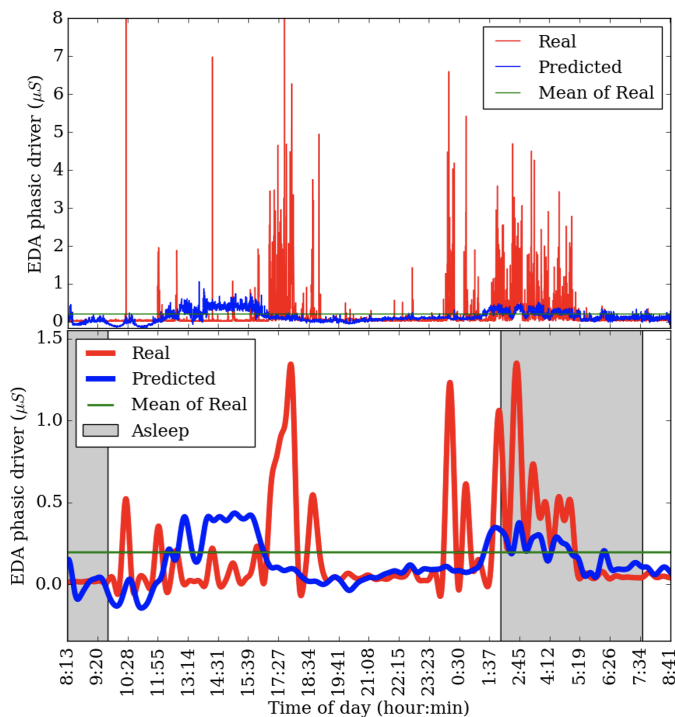


Fig. 9. Predicted vs. actual phasic driver signal for patient 2. The *top* graph shows the raw signals, while the *bottom* graph contains low-pass filtered signals.




Fig. 10. Normalised absolute DeepLIFT attribution map averaged across all SEEG spectrogram inputs when predicting the EDA for patient 1. A baseline of 0 is used for all inputs.




Fig. 11. R^2 score obtained by the network when retrained on specific contact points located in different brain regions for predicting the phasic EDA signal in patient 1.

might be the reason for the lower R^2 score of the models based on contact points in it. We can also see that signals from the white matter, covering a large region of the brain, are predictive of the sweat response.

V. DISCUSSION

This analysis is one of the first attempts to study the relationship between cerebral and peripheral physiological signals in a purely data-driven fashion leveraging the flexibility of deep neural networks. Following a supervised learning approach with the aid of neural network explainability techniques, we are able to capture and elucidate some significant relationships between peripheral and SEEG signals.

A. Findings

In most patients, our model is able to predict instantaneous heart rate changes fairly well. Some of its predictive power is likely related to the wakefulness status, *i.e.* the model is able to determine a decrease in heart rate induced by sleep based on the SEEG signals. This clearly agrees with the previous work demonstrating the efficacy of neural networks at sleep stage scoring [46], [24]. While this confounder cannot be ruled out completely, we observe that in most patients, the model captures patterns that are more complex than mere wakefulness status. We also find that the patients with contact points in the insula’s CAN region feature a considerably stronger association between the heart rate and SEEG than the patients lacking contact points in this area. Particularly, for patient 2, who had all of the electrodes placed in areas that are further away from the CAN region, no significant relationship was found between the brainwaves and both of the peripheral signals. This finding aligns well with the prior research [4], [67], [68], [49], [50]. Nevertheless, the model can still predict the heart rate using the non-CAN region contact points in other subjects. However, these contacts are in close proximity of the insula, and might still pick up signal from that area – this might be the reason for the relatively good performance, shown in Figure 7. Furthermore, using DeepLIFT attribution, we show that the heart rate is associated with the theta (4-8 Hz), alpha (8-13 Hz), and lower beta (13-20 Hz) frequency ranges, as

well as the higher frequencies (53–60 Hz). Frequencies outside these ranges are less relevant to the predictive model.

Another peripheral signal we explored in this analysis was the sweat response measured by the EDA signal’s phasic driver. The EDA is visibly less affected by the wakefulness status of the patient. Nevertheless, the model is able to predict this signal well on patients who had the contact points in the insula. Similarly to the heart rate, this resonates with previous findings described in [52], [4], [53], [68], [69], [67], [51]. For the EDA, the lower frequency bands are useful for prediction, especially the theta (4–8 Hz) and alpha (8–13 Hz) waves. This relationship is especially strong for patient 1. We also observe an apparent time delay in the attribution maps, suggesting a considerable delay, of approx. 12 to 14 seconds, between the SEEG signals and the sweat response recorded from the hand. This delay should be investigated further.

The patients selected for this study were both male and female aged 18–40 years with BMI between 20 and 40 kg/m² (see Table I). We did not observe a significant correlation between any of the mentioned patient characteristics and the strength of association between cerebral and peripheral biosignals. Moreover, we do not expect these characteristics to be confounders, since the analysis was conducted by training *patient-specific* models.

These main findings of our analysis are summarised in Table IV.

Biosignal	Heart Rate	Sweat Response
Strongest brain region relationship	CAN	CAN
Frequencies of interest	Alpha, Theta and 52–58 Hz	Alpha, Theta and 52–58 Hz
Temporal dependency	invariant	12–14 seconds delay between neuron firing and sweat glands burst

TABLE IV

SUMMARY OF FINDINGS ABOUT BRAIN-PERIPHERAL RELATIONSHIPS.

B. Limitations

This pilot study is a preliminary exploratory analysis which only includes four patients. Both a larger sample size and external validation are necessary to confirm our findings. The current analysis focuses on rather large brain regions; whereas it would be interesting to further localise relationships between the brain and peripheral biosignals by performing per contact point analysis within the the insula region.

VI. CONCLUSION

To the best of our knowledge, this is the first study correlating cerebral with peripheral biosignals applying deep neural network models, with the capacity of representing high-dimensional and nonlinear relations. In contrast, earlier studies applied simplified mechanistic models supported by experiments or utilized multivariate linear regression models. Furthermore, we have demonstrated the applicability of explainability methods, enabling a data-driven approach in the

discovery of novel patterns, governing the relation between the brain and peripheral biosignals.

The proposed analysis pipeline consists of (i) signal pre-processing to yield spectrograms from SEEG recordings and ground-truth heart rate and EDA target values, (ii) the CNN to learn relations between SEEG and biosignals, (iii) a gradient-based explanation method (i.e. DeepLIFT) to identify most relevant frequency and temporal features, and (iv) remove-and-retrain explanations to explore the correlation with specific brain regions.

The application of the analysis pipeline resulted in the following findings: strongest brain-peripheral correlations were observed from contacts located in the CAN region of the brain, for the alpha, theta and 52 to 58 Hz frequency band. A temporal delay of 12 to 14 s between SEEG and EDA signal was observed. These findings resulted from patient-specific models, required due to the unique placement of electrodes and contact points in each of the patients brain.

In the future, data from a larger patient cohort, monitored over a longer time period, including a larger variety of activities could be considered to identify more subtle and generalized patterns between brain and peripheral signals. Sequential models could be another promising extension of the current analysis, for example, an RNN or LSTM with convolutional layers could be more suitable for exploring temporal dependencies in the data. In addition, transfer learning could be leveraged to pretrain and refine the CNN on cohort and individuals data, respectively. Furthermore, the explainability methods could be applied down to a spatial resolution of single contact points. Finally, not only the correlation, but also the causality between the signals could be explored *e.g.* in a Granger-causal framework.

REFERENCES

- [1] O. D. Lara and M. A. Labrador, “A survey on human activity recognition using wearable sensors,” *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1192–1209, 2013.
- [2] V. Stanford, “Biosignals offer potential for direct interfaces and health monitoring,” *IEEE Pervasive Computing*, vol. 3, no. 1, pp. 99–103, 2004.
- [3] M. Fredrikson, T. Furmark, M. T. Olsson, H. Fischer, J. Andersson, and B. Långström, “Functional neuroanatomical correlates of electrodermal activity: a positron emission tomographic study,” *Psychophysiology*, vol. 35, no. 2, pp. 179–185, 1998.
- [4] E. E. Benarroch, “The central autonomic network: functional organization, dysfunction, and perspective,” in *Mayo Clinic Proceedings*, vol. 68, no. 10. Elsevier, 1993, pp. 988–1001.
- [5] J. F. Thayer and R. D. Lane, “A model of neurovisceral integration in emotion regulation and dysregulation,” *Journal of affective disorders*, vol. 61, no. 3, pp. 201–216, 2000.
- [6] Y. LeCun, Y. Bengio *et al.*, “Convolutional networks for images, speech, and time series,” *The handbook of brain theory and neural networks*, vol. 3361, no. 10, p. 1995, 1995.
- [7] E. Inc., “Empatica e4,” 2020. [Online]. Available: <https://www.empatica.com/en-int/research/e4/>
- [8] M. Ancona, E. Ceolini, C. Öztireli, and M. Gross, “Towards better understanding of gradient-based attribution methods for deep neural networks,” *arXiv preprint arXiv:1711.06104*, 2017.
- [9] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural networks*, vol. 61, pp. 85–117, 2015.
- [10] S. Bhanja and A. Das, “Deep neural network for multivariate time-series forecasting,” in *Proceedings of International Conference on Frontiers in Computing and Systems*. Springer, 2021, pp. 267–277.

- [11] A. Tealab, "Time series forecasting using artificial neural networks methodologies: A systematic review," *Future Computing and Informatics Journal*, vol. 3, no. 2, pp. 334–340, 2018.
- [12] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [13] C. Lea, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks: A unified approach to action segmentation," in *European Conference on Computer Vision*. Springer, 2016, pp. 47–54.
- [14] N. K. Ahmed, A. F. Atiya, N. E. Gayar, and H. El-Shishiny, "An empirical comparison of machine learning models for time series forecasting," *Econometric Reviews*, vol. 29, no. 5-6, pp. 594–621, 2010.
- [15] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016, <http://www.deeplearningbook.org>. "Last accessed 08 December 2020".
- [16] H. Lütkepohl, *New Introduction to Multiple Time Series Analysis*. Springer, 2007.
- [17] A. Shrikumar, P. Greenside, and A. Kundaje, "Learning important features through propagating activation differences," *arXiv preprint arXiv:1704.02685*, 2017.
- [18] R. Marcinkevičius and J. E. Vogt, "Interpretability and explainability: A machine learning zoo mini-tour." 2020, arXiv:2012.01805.
- [19] M. D. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *European conference on computer vision*. Springer, 2014, pp. 818–833.
- [20] J. Castro, D. Gómez, and J. Tejada, "Polynomial calculation of the Shapley value based on sampling," *Computers & Operations Research*, vol. 36, no. 5, pp. 1726 – 1730, 2009, selected papers presented at the Tenth International Symposium on Locational Decisions (ISOLDE X).
- [21] A. Fisher, C. Rudin, and F. Dominici, "All models are wrong, but many are useful: Learning a variable's importance by studying an entire class of prediction models simultaneously," *Journal of Machine Learning Research*, vol. 20, no. 177, pp. 1–81, 2019.
- [22] R. Assaf, I. Giurgiu, F. Bagehorn, and A. Schumann, "MTEX-CNN: Multivariate time series explanations for predictions with convolutional neural networks," in *2019 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2019, pp. 952–957.
- [23] A. Craik, Y. He, and J. L. Contreras-Vidal, "Deep learning for electroencephalogram (EEG) classification tasks: a review," *Journal of neural engineering*, vol. 16, no. 3, p. 031001, 2019.
- [24] D. Miladinović, C. Muheim, S. Bauer, A. Spinnler, D. Noain, M. Bandarabadi, B. Gallusser, G. Krummhacher, C. Baumann, A. Adamantidis *et al.*, "SPINDLE: End-to-end learning from EEG/EMG to extrapolate animal sleep scoring across experimental settings, labs and species," *PLoS computational biology*, vol. 15, no. 4, p. e1006968, 2019.
- [25] A. Antoniadou, L. Spyrou, C. C. Took, and S. Sanei, "Deep learning for epileptic intracranial EEG data," in *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, 2016, pp. 1–6.
- [26] A. Antoniadou, L. Spyrou, D. Martin-Lopez, A. Valentin, G. Alarcon, S. Sanei, and C. C. Took, "Detection of interictal discharges with convolutional neural networks using discrete ordered multichannel intracranial EEG," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 12, pp. 2285–2294, 2017.
- [27] S. M. Usman, S. Khalid, and M. H. Aslam, "Epileptic seizures prediction using deep learning techniques," *Ieee Access*, vol. 8, pp. 39 998–40 007, 2020.
- [28] M. Abou Jaoude, J. Jing, H. Sun, C. S. Jacobs, K. R. Pellerin, M. B. Westover, S. S. Cash, and A. D. Lam, "Detection of mesial temporal lobe epileptiform discharges on intracranial electrodes using deep learning," *Clinical Neurophysiology*, vol. 131, no. 1, pp. 133–141, 2020.
- [29] T. Song, W. Zheng, P. Song, and Z. Cui, "EEG emotion recognition using dynamical graph convolutional neural networks," *IEEE Transactions on Affective Computing*, vol. 11, no. 3, pp. 532–541, 2018.
- [30] F. Wang, S.-h. Zhong, J. Peng, J. Jiang, and Y. Liu, "Data augmentation for EEG-based emotion recognition with deep convolutional neural networks," in *International Conference on Multimedia Modeling*. Springer, 2018, pp. 82–93.
- [31] J. Chen, P. Zhang, Z. Mao, Y. Huang, D. Jiang, and Y. Zhang, "Accurate EEG-based emotion recognition on combined features using deep convolutional neural networks," *IEEE Access*, vol. 7, pp. 44 317–44 328, 2019.
- [32] P. Zhang, X. Wang, W. Zhang, and J. Chen, "Learning spatial-spectral-temporal EEG features with recurrent 3D convolutional neural networks for cross-task mental workload assessment," *IEEE Transactions on neural systems and rehabilitation engineering*, vol. 27, no. 1, pp. 31–42, 2018.
- [33] D. D. Chakladar, S. Dey, P. P. Roy, and D. P. Dogra, "EEG-based mental workload estimation using deep BLSTM-LSTM network and evolutionary algorithm," *Biomedical Signal Processing and Control*, vol. 60, p. 101989, 2020.
- [34] Y. Zhao and L. He, "Deep learning in the EEG diagnosis of Alzheimer's disease," in *Asian conference on computer vision*. Springer, 2014, pp. 340–353.
- [35] Y. Guo, K. Friston, F. Aldo, S. Hill, and H. Peng, *Brain Informatics and Health: 8th International Conference, BIH 2015, London, UK, August 30-September 2, 2015. Proceedings*. Springer, 2015, vol. 9250.
- [36] U. R. Acharya, S. L. Oh, Y. Hagiwara, J. H. Tan, H. Adeli, and D. P. Subha, "Automated EEG-based screening of depression using deep convolutional neural network," *Computer methods and programs in biomedicine*, vol. 161, pp. 103–113, 2018.
- [37] A. Supratak, H. Dong, C. Wu, and Y. Guo, "DeepSleepNet: A model for automatic sleep stage scoring based on raw single-channel EEG," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 25, no. 11, pp. 1998–2008, 2017.
- [38] Z. Yu and J. Song, "Multi-class motor imagery classification by singular value decomposition and deep Boltzmann machine," in *2017 IEEE 3rd Information Technology and Mechatronics Engineering Conference (ITOEC)*. IEEE, 2017, pp. 376–379.
- [39] Z. Tang, C. Li, and S. Sun, "Single-trial EEG classification of motor imagery using deep convolutional neural networks," *Optik*, vol. 130, pp. 11–18, 2017.
- [40] Y. R. Tabar and U. Halici, "A novel deep learning approach for classification of EEG motor imagery signals," *Journal of neural engineering*, vol. 14, no. 1, p. 016003, 2016.
- [41] X. Wei, L. Zhou, Z. Chen, L. Zhang, and Y. Zhou, "Automatic seizure detection using three-dimensional CNN based on multi-channel EEG," *BMC medical informatics and decision making*, vol. 18, no. 5, pp. 71–80, 2018.
- [42] R. J. Kobler and R. Scherer, "Restricted Boltzmann machines in sensory motor rhythm brain-computer interfacing: A study on inter-subject transfer and co-adaptation," in *2016 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2016, pp. 000 469–000 474.
- [43] L. Vidyaratne, A. Glandon, M. Alam, and K. M. Iftekharuddin, "Deep recurrent neural network for seizure detection," in *2016 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2016, pp. 1202–1207.
- [44] Z. Yin and J. Zhang, "Cross-session classification of mental workload levels using EEG and an adaptive deep learning model," *Biomedical Signal Processing and Control*, vol. 33, pp. 30–47, 2017.
- [45] Q. She, B. Hu, Z. Luo, T. Nguyen, and Y. Zhang, "A hierarchical semi-supervised extreme learning machine method for EEG recognition," *Medical & biological engineering & computing*, vol. 57, no. 1, pp. 147–157, 2019.
- [46] L. Yuan and J. Cao, "Patients' EEG data analysis via spectrogram image with a convolution neural network," in *International Conference on Intelligent Decision Technologies*. Springer, 2017, pp. 13–21.
- [47] Y. Li, J. Huang, H. Zhou, and N. Zhong, "Human emotion recognition with electroencephalographic multidimensional features by hybrid deep neural networks," *Applied Sciences*, vol. 7, no. 10, p. 1060, 2017.
- [48] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, pp. 1097–1105, 2012.
- [49] S. Oppenheimer and D. Cechetto, *The Insular Cortex and the Regulation of Cardiac Function*. American Cancer Society, 2016, pp. 1081–1133.
- [50] F. Chouchou, F. Mauguière, O. Vallayer, H. Catenoux, J. Isnard, A. Montavont, J. Jung, V. Pichot, S. Rheims, and L. Mazzola, "How the insula speaks to the heart: Cardiac responses to insular stimulation in humans," *Human Brain Mapping*, vol. 40, no. 9, pp. 2611–2622, 2019.
- [51] M. J. Farrell, D. Trevaks, N. A. Taylor, and R. M. McAllen, "Regional brain responses associated with thermogenic and psychogenic sweating events in humans," *Journal of neurophysiology*, vol. 114, no. 5, pp. 2578–2587, 2015.
- [52] M. E. Dawson, A. M. Schell, and D. L. Fillion, "The electrodermal system," in *Principles of Psychophysiology: Physical, Social, and Inferential Elements*. Cambridge University Press, 1990, pp. 295–324.
- [53] H. D. Critchley, "Review: Electrodermal responses: What happens in the brain," *The Neuroscientist*, vol. 8, no. 2, pp. 132–142, 2002.

- [54] H. van Lier, M. Pieterse, A. Garde, M. Postel, H. de Haan, M. Vollebrouk Hutten, J. M. Schraagen, and M. Noordzij, "A standardized validity assessment protocol for physiological signals from wearable technology: Methodological underpinnings and an application to the e4 biosensor," *Behavior Research Methods*, vol. 52, 07 2019.
- [55] S.-P. Kim, "Preprocessing of EEG," in *Computational EEG Analysis*. Springer, 2018, pp. 15–33.
- [56] R. Lyons, "Understanding cascaded integrator-comb filters," *Embed Syst Program*, vol. 18, no. 4, pp. 14–27, 2005.
- [57] B. Rajoub, "Characterization of biomedical signals: Feature engineering and extraction," in *Biomedical Signal Processing and Artificial Intelligence in Healthcare*. Elsevier, 2020, pp. 29–50.
- [58] M. Benedek and C. Kaernbach, "A continuous measure of phasic electrodermal activity," *Journal of neuroscience methods*, vol. 190, no. 1, pp. 80–91, 2010.
- [59] —, "Decomposition of skin conductance data by means of nonnegative deconvolution," *Psychophysiology*, vol. 47, no. 4, pp. 647–658, 2010.
- [60] L. Wyse, "Audio spectrogram representations for processing with convolutional neural networks," *arXiv preprint arXiv:1706.09559*, 2017.
- [61] R. Hyder, S. Ghaffarzadegan, Z. Feng, J. H. Hansen, and T. Hasan, "Acoustic scene classification using a CNN-SuperVector system trained with auditory and spectrogram image features," in *Interspeech*, 2017, pp. 3073–3077.
- [62] C. Ito, X. Cao, M. Shuzo, and E. Maeda, "Application of CNN for human activity recognition with FFT spectrogram of acceleration and gyro sensors," in *Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers*, 2018, pp. 1503–1510.
- [63] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035.
- [64] D. R. Roberts, V. Bahn, S. Ciuti, M. S. Boyce, J. Elith, G. Guillera-Arroita, S. Hauenstein, J. J. Lahoz-Monfort, B. Schröder, W. Thuiller, D. I. Warton, B. A. Wintle, F. Hartig, and C. F. Dormann, "Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure," *Ecography*, vol. 40, no. 8, pp. 913–929, 2017.
- [65] D. Anguita, L. Ghelardoni, A. Ghio, L. Oneto, and S. Ridella, "The 'k' in k-fold cross validation," in *ESANN*, 2012, pp. 441–446.
- [66] S. Wright, "The theory of path coefficients a reply to Niles's criticism," *Genetics*, vol. 8, no. 3, p. 239, 1923.
- [67] J. K. Shoemaker, K. N. Norton, J. Baker, and T. Luchyshyn, "Forebrain organization for autonomic cardiovascular control," *Autonomic Neuroscience*, vol. 188, pp. 5–9, 2015.
- [68] J. F. Thayer, F. Åhs, M. Fredrikson, J. J. Sollers III, and T. D. Wager, "A meta-analysis of heart rate variability and neuroimaging studies: implications for heart rate variability as a marker of stress and health," *Neuroscience & Biobehavioral Reviews*, vol. 36, no. 2, pp. 747–756, 2012.
- [69] M. J. Farrell, D. Trevaks, N. A. Taylor, and R. M. McAllen, "Brain stem representation of thermal and psychogenic sweating in humans," *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, vol. 304, no. 10, pp. R810–R817, 2013.