

Removing the cardiac field artifact from the EEG using neural network regression

Stefan Arnau¹  | Fariba Sharifian² | Edmund Wascher¹  | Mauro F. Larra¹ 

¹Leibniz Research Centre for Working Environment and Human Factors Dortmund (IfAdo), Dortmund, Germany

²School of Computer Science and Mathematics, Liverpool John Moores University, Liverpool, UK

Correspondence

Stefan Arnau, Leibniz Research Centre for Working Environment and Human Factors Dortmund (IfAdo), Ardeystraße 67, 44139 Dortmund, Germany.

Email: arnau@ifado.de

Funding information

German Research Foundation, Grant/Award Number: DFG LA 4830/2-1

Abstract

When EEG recordings are used to reveal interactions between central-nervous and cardiovascular processes, the cardiac field artifact (CFA) poses a major challenge. Because the electric field generated by cardiac activity is also captured by scalp electrodes, the CFA arises as a heavy contaminant whenever EEG data are analyzed time-locked to cardio-electric events. A typical example is measuring stimulus-evoked potentials elicited at different phases of the cardiac cycle. Here, we present a nonlinear regression method deploying neural networks that allows to remove the CFA from the EEG signal in such scenarios. We train neural network models to predict R-peak centered EEG episodes based on the ECG and additional CFA-related information. In a second step, these trained models are used to predict and consequently remove the CFA in EEG episodes containing visual stimulation occurring time-locked to the ECG. We show that removing these predictions from the signal effectively removes the CFA without affecting the intertrial phase coherence of stimulus-evoked activity. In addition, we provide the results of an extensive grid search suggesting a set of appropriate model hyperparameters. The proposed method offers a replicable way of removing the CFA on the single-trial level, without affecting stimulus-related variance occurring time-locked to cardiac events. Disentangling the cardiac field artifact (CFA) from the EEG signal is a major challenge when investigating the neurocognitive impact of cardioafferent traffic by means of the EEG. When stimuli are presented time-locked to the cardiac cycle, both sources of variance are systematically confounded. Here, we propose a regression-based approach deploying neural network models to remove the CFA from the EEG. This approach effectively removes the CFA on a single-trial level and is purely data-driven, providing replicable results.

KEYWORDS

artifact removal, cardiac field artifact, ECG, EEG, neural networks, regression

1 | INTRODUCTION

The muscular activity of the heart creates an electric field that propagates throughout the entire body (Dirlich et al., 1997). The electric field of the heart is 1–2 mV in amplitude and typically quantified via ECG, but is also captured by the electrodes of the EEG, producing a regularly occurring artifact, the cardiac field artifact (CFA). In experimental paradigms commonly used in EEG research, CFA-related variance is usually unsystematic and therefore diminished by averaging procedures when parameterizing the signal. Still, the representation of the CFA in the EEG signal might reduce the signal-to-noise ratio (SNR) of the phenomenon of interest. In experiments investigating the central-nervous effects of cardioafferent traffic, however, these phenomena happen to be time-locked to the activity of the heart. Specifically, to study how phasic variations in cardioafferent traffic impact information processing (Azzalini et al., 2019; Critchley & Harrison, 2013; Garfinkel & Critchley, 2016), a standard paradigm is to present brief stimuli at times of low and high cardioafferent traffic by synchronizing their onset to the electrocardiogram (ECG; Al et al., 2020; Azevedo et al., 2017; Garfinkel et al., 2014; Gray et al., 2010; Larra et al., 2020; Pramme et al., 2016; Schulz et al., 2009). Consequently, EEG signal and CFA-related variance are confounded systematically, and removing the CFA is necessary to interpret the electrophysiological findings. Thus, a major challenge when investigating the neurocognitive impact of cardioafferent traffic by means of the EEG is a systematic confound between variance induced by cardioafferent traffic and variance produced by the electrical field of the heart.

Several methods may be used to suppress the CFA in the recorded EEG signal. Some studies used specific derivations or transformations to attenuate the CFA (Hjorth, 1975; Luft & Bhattacharya, 2015; Terhaar et al., 2012). Another approach is to correct for the CFA by subtracting a baseline ERP without stimulation from the ERP of experimental trials (Gray et al., 2010; van Elk et al., 2014). Most recent studies rely on signal decomposition using independent component analyses (ICA; Hyvärinen et al., 2001), a blind source separation method that aims at decomposing a mixed signal into its statistical sources (c.f. Ullsperger & Debener, 2010). ICA can be used to detect and remove artifacts (Jung et al., 2000) by identifying and removing those independent components (ICs) that represent mostly artifactual activity. Identifying ICs reflecting cardiac activity can be done manually, based on topography and time course (Al et al., 2020, 2021; Bury et al., 2019; Ullsperger & Debener, 2010), or with the help of classification tools (Issa et al., 2019; Pion-Tonachini et al., 2019; Radüntz et al., 2017; Tamburro et al., 2019;

Viola et al., 2009). Despite such classification tools, the decision of whether a given IC represents primarily cardiac or brain activity is often difficult and not always objectifiable, thus not replicable.

Since the ECG source signal is readily available, a regression-based technique would seem to be another obvious method of choice. This possibility, however, has been left surprisingly unexplored. Here, we present a regression-based approach to remove the CFA from the EEG signal when the experimental stimulation is time-locked to the ECG. An important consideration when designing regression models to predict the CFA from the ECG signal is that the relationship may not be entirely linear but will partly depend on the type of ECG derivation as well as on the cardiac cycle. Over the course of the cardiac cycle, the electrically active cardiac structures change. Therefore, the dipole generating the CFA changes its spatial orientation relative to the EEG montage, inducing temporally specific dependencies between the EEG and ECG signals. To address this, we entered this information as predictors and used neural network models with dense hidden layers to perform the regression task, as these models are capable to model such nonlinear dependencies. Recent studies successfully deployed neural network models for cleaning the EEG from various types of artifacts (Lopes et al., 2021; Merlin Praveena et al., 2022). An advantage of neural network models is that they are easily scalable in terms of their complexity. It should thus be possible to define models that are sufficiently complex to perform well, while also keeping computational demands to a minimum.

For the present study, we recorded the EEG as well as the ECG of 40 participants during an experiment in which lateralized visual stimuli were presented time-locked to the ECG using a closed-loop system. Our goal was to remove the CFA on a single trial level with maximum efficacy while preserving specific EEG variance related to the visual stimuli. First, we present the results of an extensive grid search conducted to determine which set of model hyperparameters provides an optimal balance of speed and performance. These results are presented to allow future studies to easily employ this method, since performing a grid search to fine-tune the model hyperparameters is very computation-intensive and time-consuming. Subsequently, we trained the neural networks to predict the EEG from the ECG (and additional features), and then removed the CFA by subtracting these predictions from the signal. We compared the ERPs of the contaminated and the cleaned data by calculating effect sizes for each point in time \times sensor space as an estimate of the efficacy of the CFA removal approach. Finally, we calculated the inter-trial phase coherences (ITPCs) for the contaminated and the cleaned

signals to demonstrate that this approach is appropriate to remove the CFA on the single trial level while preserving stimulus-related information. Here, effect sizes for time \times frequency \times sensor data are provided.

2 | METHOD

2.1 | CFA removal using neural network regression

The basic idea of the method for removing the CFA from the EEG signal presented here is to use a regression modeling approach. Regression model training is performed on EEG data without systematic experimental stimulation. The models are trained for each participant and electrode to predict the EEG signal using the ECG signal and temporal features of the cardiac cycle. The data used for training the models consisted of individual time points from epochs centered at R-peaks that neither preceded nor succeeded an experimental stimulation event. For evaluating the models' performance, 20% of the samples were put aside as a test dataset and were not used for training. During training, 20% of the remaining samples were also not used for training, but as a validation-dataset to track the model fit over training iterations. After training, predictions were made for the experimental dataset, containing samples from episodes of stimulus-induced activity time-locked to the ECG. These predictions then get subtracted from the observed signal for each EEG channel and participant individually, so that anatomical differences, differences between the recording-sessions and the spatial orientations of the dipoles generating the CFA relative to the EEG montage are considered. This also means that the method proposed here should produce comparable results independent of population characteristics as sources of interindividual variability are taken into account. It is therefore not advisable to rely on pretrained models (i.e., from a different sample) as this could be detrimental to performance.

We deployed neural network models composed of dense layers with an equal number of neurons using Tensorflow (Abadi et al., 2015) with the Keras API (Chollet et al., 2015) for Python. An important design choice when specifying neural network architectures is to define the model's hyperparameters. Some choices regarding hyperparameters were based on best practice recommendations: We used the ELU activation function (Clevert et al., 2016), as it was shown that it does suffer neither from vanishing or exploding gradients, nor from dying neurons. It has also been described as converging comparably fast during training. For weight initialization, we used the method proposed by He

et al. (2015) and we used Nadam as an optimizer for the gradient descent algorithm (Dozat, 2016). For regularization, to prevent overfitting the training data, we made use of early stopping during training. Training stopped when the loss metric did not further improve over the course of ten epochs. The mean squared error was used as the loss function, and the loss was calculated based on validation data that was not used for training.

Beside these a priori design choices, we also performed a grid search for several other hyperparameters and model features to investigate the models' performance in terms of prediction accuracy and training speed. One pair of hyperparameters determining the complexity of the model are the number of layers and the number of neurons per layer. Our grid search encompassed 1, 2, 4, or 8-layer models consisting of 5, 10, 20, or 50 neurons each. A hyperparameter choice possibly affecting model accuracy and training speed is batch size, which determines the amount of training data (i.e., samples/time-point data) the model uses for each iteration. Here, we compare a smaller batch size of 128 to a larger batch size of 2048. Similarly, we investigated how far a reduction of the training data affects the models' performance by using either all available training data or just 25% of it.

A possible source of nonlinearity in the relation between ECG recorded at the chest and CFA recorded at the scalp are changes in the location and orientation of the dipole generating the electric field that produces the CFA. These changes should be related to the cardiac cycle. In the grid search, we therefore also compared whether information on the cardiac cycle helps to improve the model. To this end, we created two additional features (time series). One of these features simply consisted of the time in milliseconds that had passed since the last R-peak. The second feature consisted of the percentage of time that had already passed in the interval from the last to the next R-peak. We also investigated which ECG electrode's signal to use for the model. In our setup, we used the ECG electrodes of the cardiac trigger monitor as well as a dropdown electrode from the EEG cap, which we attached to the chest. Since this dropdown electrode is referenced against a scalp electrode, its signal should be sensitive to head movements and thus provide additional information. In the grid search, we compared the models' performance when using either one or both of these electrodes' signals as features. Overall, this selection leads to $4 \text{ (layers)} * 4 \text{ (neurons)} * 2 \text{ (batch size)} * 2 \text{ (training dataset size)} * 2 \text{ (RR-interval latency and phase as additional features versus no information)} * 3 \text{ (ECG channels)} = 384$ models to train during the grid search to estimate the optimal hyperparameter set for the intended task. The results of the grid search are reported in the section "Feature selection and hyperparameter tuning."

2.2 | Sample

The sample consisted of $N=40$ (20 female) healthy, right-handed participants, aged 19 to 31 years ($M=23.78$, $SD=3.42$) with (corrected to) normal vision. Applicants were not included if they showed any evidence of acute or chronic diseases of the cardiovascular system (deviations from sine rhythm, glaucoma, history of fainting, resting blood pressure above 140/90 mmHg), history of psychiatric disease or neurological disorders. All exclusion and inclusion criteria were assessed during a brief telephone interview. Normal blood pressure and sine rhythm (i.e., absence of extrasystoles) were confirmed during a ten-minute resting phase before the start of the experiment. Further exclusion criteria were smoking of more than five cigarettes per day, drug intake, or current use of medication. Participants were recruited by announcement in the institute's social media accounts. The study was performed in accordance with the Declaration of Helsinki. All participants took part voluntarily and gave written informed consent. The ethics committee of the Leibniz Research Centre for Working Environment and Human Factors, Dortmund approved the study.

2.3 | Experimental procedure

All participants were invited to the laboratory at 8:30 a.m. After the preparation of the EEG cap and the ECG

electrodes, the participants were seated in a sound attenuated EEG recording cabin with dim lighting. The participants were then fitted with a medical cervical collar with the intention of restricting head movements as much as possible. During the experiment, the letters “S” and “X” were presented as target stimuli, one at each trial (see Figure 1a), for 50 ms at a lateralized position with an offset of 2° visual angle relative to the center of the screen. To balance the physical properties of the visual stimulation, a noise stimulus consisting of three horizontal bars was presented simultaneously, contralateral to the target stimulus. The height and width of the target and noise stimuli were at a 0.45° visual angle. The lateralized positions for the stimulus presentation were outlined by squares with an edge length of 1° that were visible the whole time. Also, a fixation cross with a diameter of 0.3° was presented in the center of the screen. All visual stimuli were presented in light gray (CIE values: $x=0.287$, $y=0.312$, $Y=80$) against a dark background (CIE values: $x=0.287$, $y=0.312$, $Y=10$). The participants were instructed to identify both letters “S” and “X” by pushing a left or right button using the index finger of their respective hands. Whether the letter “S” or “X” was assigned to the left or right index finger was counterbalanced across participants. They were instructed to respond as quickly and accurately as possible. Overall, the experiment consisted of a practice block of 5 min and four experimental blocks of 20 min each (see Figure 1c).

In the four experimental blocks, the target stimuli were presented time-locked to the R-peaks of the ECG to

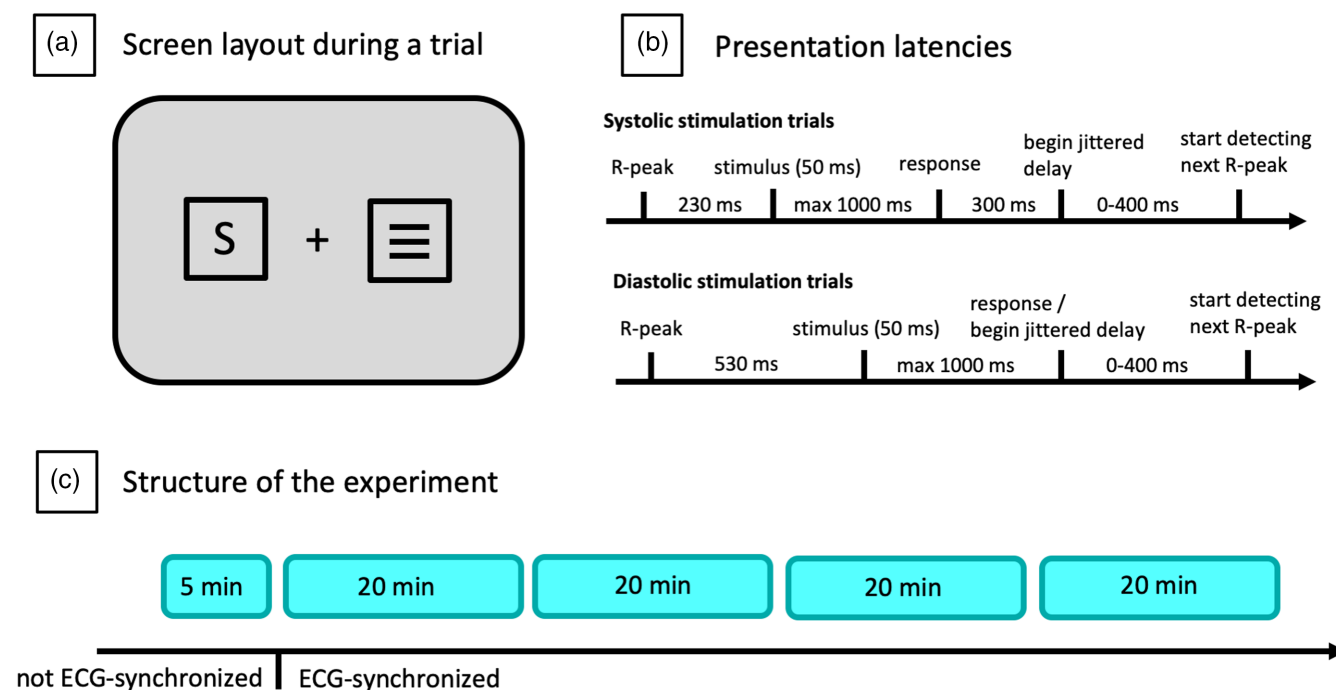


FIGURE 1 This figure depicts in (a) an example of the experimental stimuli used, in (b) the stimulus and delay period latencies with respect to the R-peak during systolic and diastolic stimulation, and in (c) an overview of the block-structure of the experiment.

stimulate either 230 ms (systole) or 530 ms (diastole) after the R-peak (Schulz et al., 2009). R-peaks were detected online using an AccuSync ECG monitor (AccuSync 72, Medical Research Corp.) which delivered TTL pulses to the stimulus presenting PC, creating a closed-loop system. After a participant responded to a stimulus or the maximum response time was reached, there was a jittery delay interval with an average length of 700 ms (systolic trials) or 400 ms (diastolic trials). The 300 ms difference in this interval was introduced to compensate for the different onset latencies between systolic and diastolic stimulation. The subsequent R-peak after this delay interval was then used as the time-locking event for the next trial. This concept is illustrated in Figure 1b. Systolic and diastolic stimulation were altered block-wise, with half of the participants beginning with systolic stimulation in the first experimental block and the other half beginning with diastolic stimulation.

2.4 | EEG and ECG acquisition

A montage of 60 active Ag/AgCl electrodes (ActiCap, Brain Products) arranged according to the 10–20 system was used to record the EEG. During recording, the Fpz electrode was used as ground, and the P9 electrode served as an online reference. ECG data were recorded using a bipolar montage of two Ag/AgCl electrodes placed in Lead II configuration. Another electrode was placed below the chest and referenced to P9 to approximate the CFA at the scalp electrodes. Data were sampled at 1000 Hz by using a BrainAmp DC amplifier (Brain Product). A hardware low-pass filter was applied at 250 Hz, and the impedances were kept below 10 k Ω .

2.5 | EEG preprocessing

The open-source toolbox EEGLab (Delorme & Makeig, 2004) was used to preprocess the EEG data with MATLAB 2019b (The Math Works Inc.). At first, EEG data as well as ECG data were band-pass filtered from 0.5 to 30 Hz using a Butterworth filter with filter order 6. With respect to the ECG data, two additional time series were created: one time series representing the cardiac cycle at each time point ranging from 0 to 1, and another time series containing the time passed since the last R-peak in sampling points. The 59 EEG channels were then scanned for electrodes with a bad signal. On average, 0.63 channels ($SD=0.81$) were excluded from further analysis. Afterward, a common average reference was used to re-reference the EEG data. For computing an Independent Component Analysis (ICA), the data were resampled at

200 Hz, epoched into consecutive segments of 2000 ms length, and epochs containing bad-quality data were identified and removed using EEGLab's automatic epoch rejection function (see code for details). After computing the ICA, the IC weights were copied to the 1000 Hz continuous data set. There, ICs representing artifacts were identified using ICLabel (Pion-Tonachini et al., 2019). ICs labeled with a probability of at least 0.3 as reflecting muscle or ocular artifacts were removed unless they were labeled with a probability of at least 0.1 to reflect cardiac activity. On average, 7.3 ICs ($SD=2.88$) were removed. Finally, previously excluded channels were interpolated.

Two datasets, one stimulus-locked and another R-peak-locked, were created. We refer to the dataset time-locked to the onset of the stimuli as the experimental dataset. The accuracy of the visual stimulation regarding the cardiac phase was ensured by discarding trials not presented within the range of 30 ms of the targeted latency (230 ms for systolic stimulation, 530 ms for diastolic stimulation). On average, 514.6 trials (22.9%) were removed due to latency imprecision. The epochs of the experimental dataset ranged from –1000 to 2000 ms relative to stimulus onset. Epochs of bad data quality were then identified in the EEG data using EEGLab's automatic epoch rejection function and removed from the EEG as well as from the ECG data. On average, 175.18 epochs were excluded ($SD=117.64$) and the number of remaining trials was 1494.18 on average ($SD=226.65$).

The dataset time-locked to R-peaks will be referred to as the training dataset. It comprised only epochs centered at R-peaks that neither preceded nor succeeded an experimental stimulus. This was done to prevent systematic variance related to the experimental stimulation from being present in this dataset. Overall, 3322.58 epochs ($SD=1403.65$) could be identified that exhibited these properties. These epochs ranged from –1000 to 1000 ms relative to the R-peaks. On average, 790.25 ($SD=555.7$) were excluded due to bad data quality, leaving a total of 2532.33 ($SD=1051.25$) for further analysis. Since the regression was performed using the data from individual time-points as samples, over 5,000,000 samples were available on average before splitting the dataset into training-, validation-, and test-sets.

2.6 | Calculating the ITPCs

ITPCs were calculated using the MNE module for Python (Gramfort et al., 2013). The time-frequency decomposition was done using Morlet-wavelet convolution. Overall, 20 linearly spaced frequencies were extracted, ranging from 2 to 20 Hz by using an also linearly spaced number of cycles ranging from 3 to 12. The ITPC is a measure of phase coherence and is calculated as the circular sum of

phases at a specific time point. Here, the ITPC is calculated across trials.

2.7 | Effect sizes

To evaluate the differences between the contaminated and the cleaned signal, effect sizes are calculated as the bias-corrected partial η^2 (Mordkoff, 2019; subsequently referred to as η^2) and are also classified as small, medium, or large according to Cohen (1992). The η^2 can be used as an estimate for the efficacy and specificity of the CFA removal approach. For the ERPs, effect sizes are reported for each point in time \times sensor space. For the ITPCs, effect sizes are reported for time \times frequency \times sensor data.

3 | RESULTS

3.1 | Feature selection and hyperparameter-tuning

To estimate an optimal set of hyperparameters and features to define the neural network architecture and

training data, a grid search was performed using a selection of values for the number of layers, the number of neurons within each layer, the batch size of each iteration during training, the size of the training dataset, whether cardiac phase information is used as a feature, and which ECG channels are used as a feature. Overall, 384 models were trained for the data of each participant for the EEG channels FCz, P7, and P8. The results, averaged across participants and EEG channels, are depicted in Figure 2.

The left side of Figure 2 illustrates the results of the grid search regarding the accuracy of the model. Depicted is the mean squared error (MSE) as a measure of how good the model fits the test data after training. The test data are a proportion (20%) of the initially available data that were not used during the actual training but used after training to evaluate the model's performance. A lower MSE indicates a better model fit. The results indicate that the performance of the neural networks was generally better with more complex models, that is, with a higher number of dense hidden layers and a higher number of neurons. More importantly, however, the model's accuracy increases with the amount of information provided by the features. The performance gets better if information from both ECG channels, the lead II ECG channel as well as the

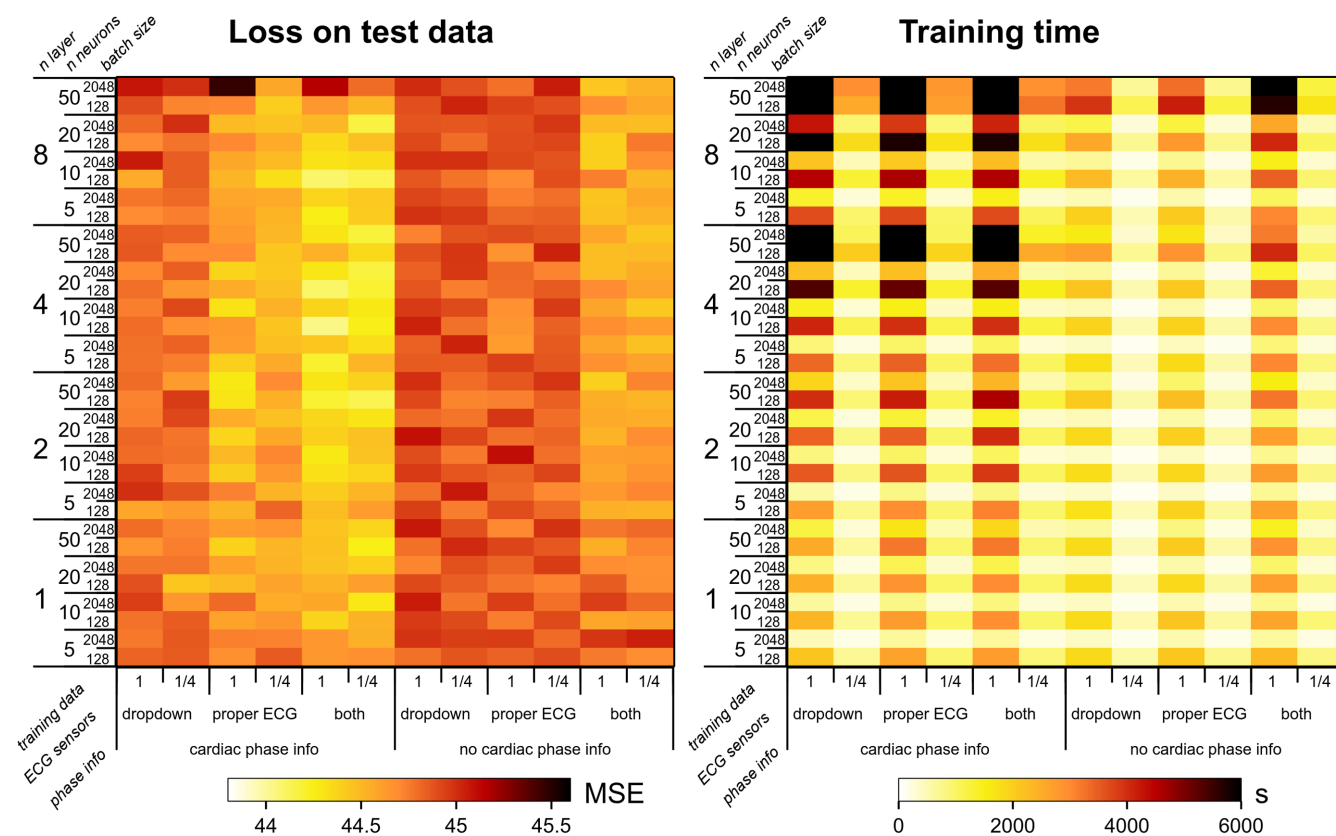


FIGURE 2 This figure depicts the results of the grid search used for hyperparameter tuning and feature selection. The left panel illustrates the model performance as the mean squared error (MSE) as a goodness of fit statistic when using the trained model on the test data. The right side shows the training time in seconds.

dropdown EEG electrode, is used. Providing information about the cardiac phase also improves performance.

The right side of Figure 2 shows the average training time for the models in seconds. Please note that the absolute values presented here are not necessarily representative, as the grid search was parallelized and each model was trained using only a single thread of the CPU. Here, the results show that training time increases with model complexity (number of layers and number of neurons within each layer) and with the number of features (number of ECG channels and cardiac phase information). Additionally, the pattern of results clearly indicates that training time is also impacted by batch size, with larger batch sizes being associated with faster training and the size of the training dataset, with a reduced training dataset leading to a faster training.

In summary, batch size and size of the training dataset affected training time substantially with comparably small effects on the model's performance, while adding information about the cardiac cycle and both ECG channels increased performance with smaller effects on training time. Thus, to balance performance and training time, it is advisable to use models of medium complexity, provide additional features, and use only a proportion of the training data.

For the following analyses, we trained the models using the data from both ECG channels as well as the cardiac phase information as features, as using more features seems to have the strongest effect on model fit. The strongest impact on training time had the factor sample set size. Therefore, we used only 25% of the available training data. With these choices made, all remaining combinations of hyperparameter settings provided an acceptable training time. We therefore went with the combination that provided the best model fit and used neural networks comprised of eight layers with ten neurons each and a batch size of 128 samples. Figure 3 illustrates the effectiveness of the regression-based CFA removal for R-peak centered epochs without experimental stimulation. The figure also illustrates that this approach reduced inter-individual variance related to the CFA without introducing inter-individual variance related to the method itself.

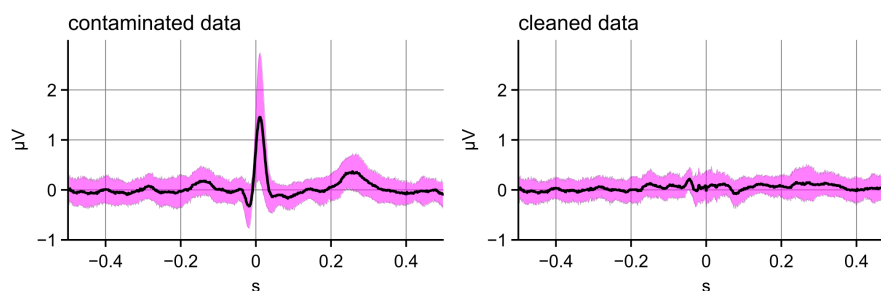


FIGURE 3 This figure shows the ERP of R-peak centered epochs without experimental stimulation for contaminated and cleaned data. The black line is the average across participants and the electrodes P7, P8, and FCz. The magenta shaded area indicates the standard deviation of the ERP across participants at each time point.

3.2 | Effect of CFA removal on ERPs

To quantify the effects of the CFA removal in ERP analyses, we compared the ERPs of the CFA-contaminated and cleaned signals of stimulus-locked epochs. The ERPs of the contaminated, predicted, and cleaned signals, as well as the associated effect sizes, are depicted in Figure 4. The effect size estimates indicate that the effects of the CFA removal are medium to large (c.f. Cohen, 1992) at left parietal, occipital, and frontal recording sites. The locations of these clusters in time suggest that the observed differences between the contaminated and the cleaned signals specifically reflect a reduction of the CFA.

3.3 | Effects of CFA removal on ITPCs

To assess the performance and specificity of the CFA removal on the single trial level, we show the effects of the CFA removal for ITPCs (see Figure 5). The effect sizes calculated for the comparison of the ITPCs obtained from the contaminated and the cleaned signals are medium to large, in particular at posterior and fronto-lateral electrodes.

4 | DISCUSSION

The goal of the present study was to investigate a regression-based approach deploying neural networks to remove the CFA from the EEG signal, which is necessary when using experimental designs in which the signal of interest is time-locked to cardiac activity. We found that such an approach represents a feasible alternative to existing methods and allows for effective removal of the CFA on a single-trial level. Moreover, we found that the performance of regression models to predict the CFA critically depends on the availability of additional information and the choice of hyperparameters.

Before the actual CFA removal, we performed a grid search to explore the performance of the model for

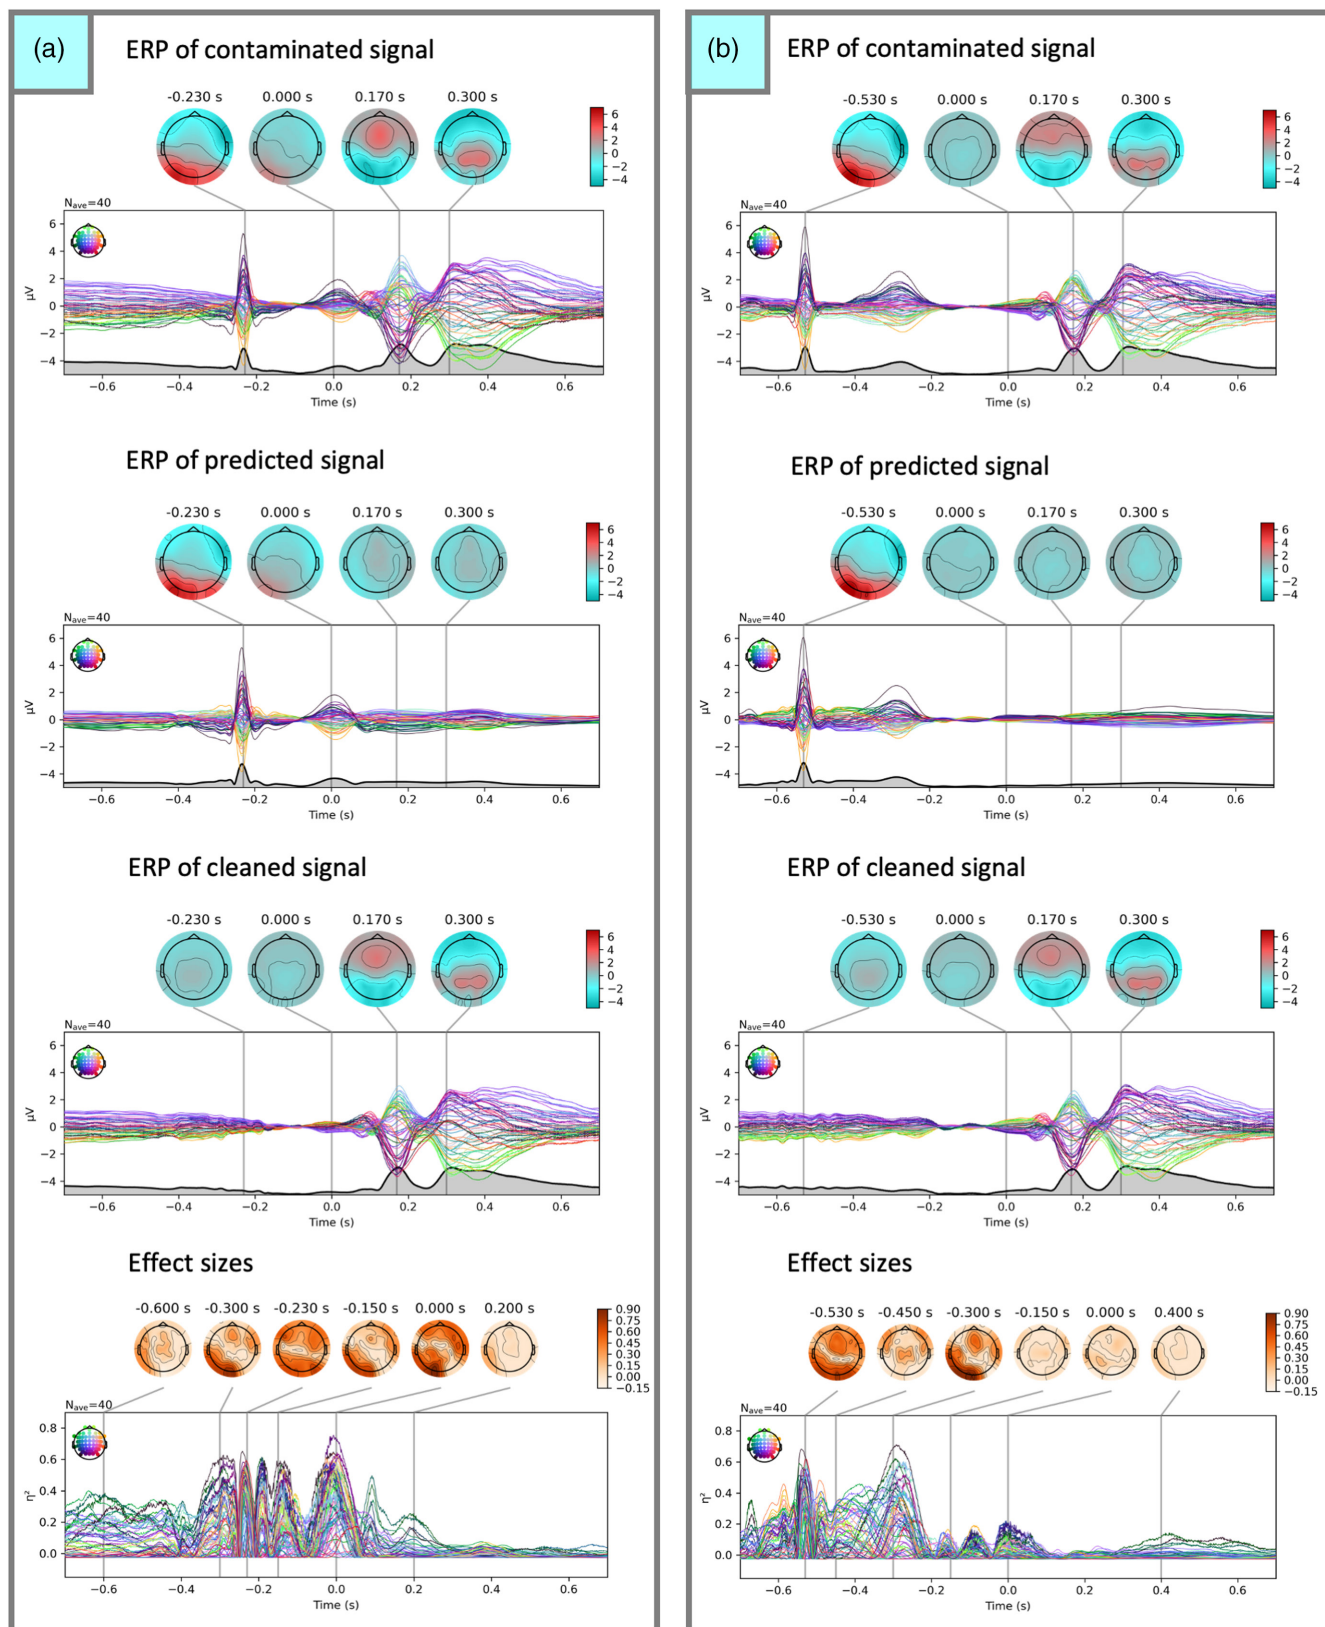


FIGURE 4 The figure shows the ERPs calculated for trials with systolic (a) and diastolic (b) experimental stimulation. Displayed are the ERPs for the contaminated, predicted, and the cleaned signal, as well as the associated effect size estimates for the comparison of the contaminated and the cleaned signals. A baseline ranging from -200 to 0 ms relative to stimulus onset was used for the ERPs. Time-point zero refers to the onset of the experimental stimuli. The black line with the gray area under the curve in the ERP plots displays the global field power.

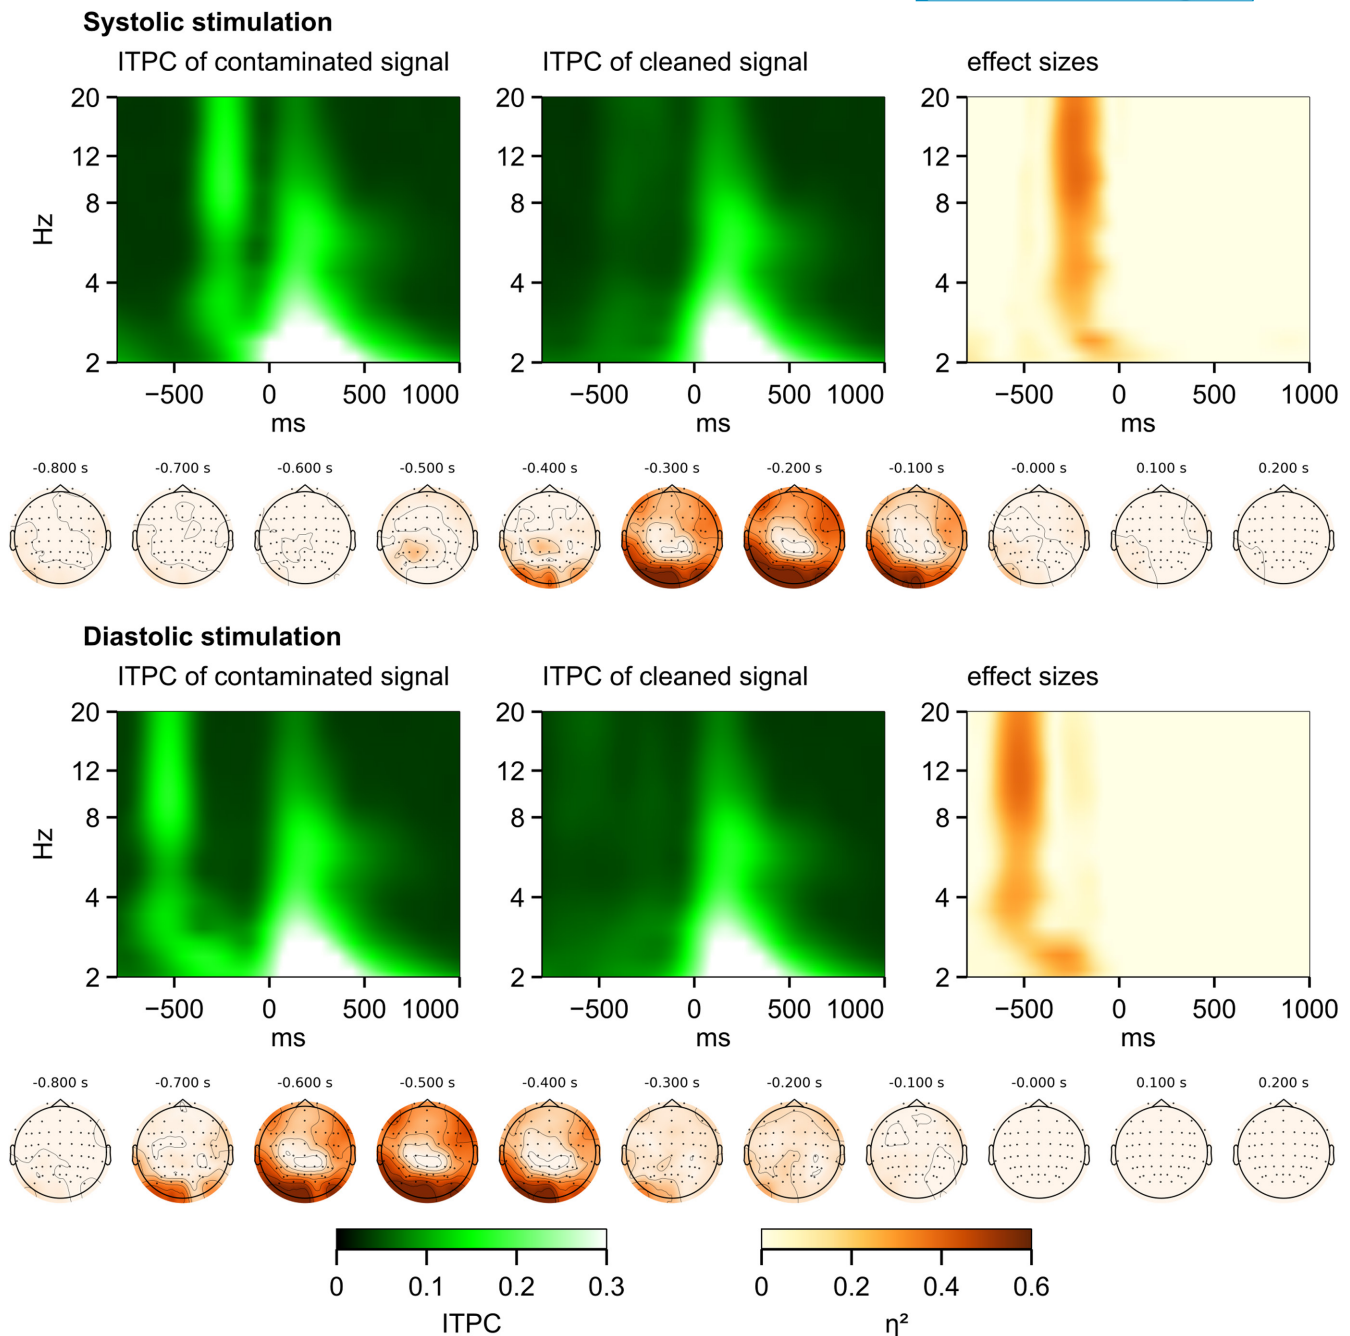


FIGURE 5 This figure shows the ITPCs for the contaminated and the cleaned signals for trials with systolic and diastolic stimulation. Effect sizes for the CFA removal are depicted in time-frequency space, averaged across channels, and as topographies for several time points throughout the trial. The frequency axis in the time-frequency plots is log-scaled.

various combinations of hyperparameter settings and features. The results provide the necessary information on how to best design such a model and, by doing so, also offer insights into the nature of the regression task and the CFA as well. Regarding the hyperparameters, the results of the grid search indicate that the models' fit improves with more complex architectures, that is, a higher number of layers and neurons, which might be necessary to model nonlinear dependencies. However, the results also indicate that the more important factor

is to provide more information when training the models. The fit also increased when both ECG channels were used as compared to only one. This could indicate that the ECG channels were nonredundant in the sense that they were sensitive to different proportions of the variance of the cardiac field. As one of the ECG channels was referenced against the head and the other against the chest, it could also be the case that the covariance of both channels provides meaningful information regarding the relative orientation of the dipole generating the

cardiac field and the EEG sensors. That is, the covariance of the ECG channels might allow for modeling head movements. Another important factor in improving the models' fit was using features that provided information about the cardiac phase. We created two channels containing cardiac phase information. One of these channels contained the time that has passed since the last R-peak; the other contained the temporal position in the current RR-interval. The observation that the models fit better when cardiac phase information was available highlights the nonlinear nature of the regression task. As the orientation and location of the dipole generating the CFA change over the course of the cardiac cycle, a given value in an ECG channel might translate to different values in the EEG channels at different cardiac phases. Eventually, we opted for using eight-layer models with ten neurons per layer. As input features, we used both ECG channels as well as information about the cardiac phase.

Another set of hyperparameters we investigated concerns the tradeoff of model fit and training time. The training of neural networks occurs in epochs. In each epoch, the model is trained on all the training data, which are presented sequentially in batches, with the models' parameters being adjusted with each batch. Smaller batch sizes thus lead to a larger number of adjustments to the models' parameters, which can lead to a better model fit. The training time of neural network depends to a large extent on the amount of training data and on how much of this data is used in each batch. Our observations suggest that much training time can be avoided when using a smaller amount of data for training and larger batch sizes without decreasing the models' fit too much.

To investigate the performance of the proposed CFA-removal approach, we compared the stimulus-locked ERPs and ITPCs of the contaminated and the cleaned signals for trials with systolic and diastolic stimulation. We chose to include the ITPC analysis to demonstrate that an analysis requiring single trial data can be performed using this approach. Overall, the results for stimulus-centered ERPs show that the approach is effective and affects the CFA specifically. In the present study, we aimed for a presentation latency relative to the R-peak of 230 ms and 530 ms for systolic and diastolic stimulation, respectively. Effect size estimates for the ERPs as well as for the ITPCs clearly show that the difference between the contaminated and the cleaned signal is the largest around these time points. For the ERPs, the effect size estimates exhibit medium to large effects (Cohen, 1992), which are most pronounced at left parietal and frontal recording sites. For the ITPCs, effect sizes were medium to large, with the strongest effects in the beta-range and at occipital electrodes. Both analyses show that the EEG signal related to cognitive processing is little affected by the CFA-removal approach.

Previously proposed approaches to remove the CFA all have specific downsides, possibly limiting their effectiveness, the reproducibility, or their range of applicability. Hjorth (1975) suggested deriving the EEG within the 10–20 system in a way that is less sensitive to signals with the spatial distribution and the special frequency the CFA exhibits (c.f. Montoya et al., 1993; Pollatos & Schandry, 2004). Based on a similar principle, transforming the data to the current source density (CSD; Kayser & Tenke, 2006) has also been proposed to attenuate the CFA by reducing its low spatial-frequency component (Luft & Bhattacharya, 2015; Terhaar et al., 2012). The downside of methods that dampen signal components with low spatial frequency is that they reduce the SNR for EEG components that exhibit such properties as well. In contrast, CFA components with a clearly defined topographical characteristic might be relatively unaffected.

Another approach focuses on temporal aspects of the CFA. In their experiment, van Elk et al. (2014) presented auditory stimuli at various latencies after the R-peak. To correct for the CFA, they subtracted a baseline ERP from the ERP of each condition. The baseline ERP was obtained from episodes without auditory stimuli. This approach has also been used in a study investigating pain-evoked potentials (Gray et al., 2010). The advantage of using ERPs time-locked to the cardiac phase as an estimator for the CFA is that it takes specific characteristics of the CFA related to the individual and recording session into account. A potential downside of this method is, that it does not work at the level of individual trials. Certain analysis methods that rely on single trials cannot be deployed. This also means that this approach cannot remove the proportion of CFA-related variance from the signal that reflects the trial to trial variability of the CFA.

The approach that is by far the most commonly used to remove the CFA in recent studies is independent component analysis (ICA; Hyvärinen et al., 2001). ICA is a blind source separation method that aims at decomposing a mixed signal into statistical sources (c.f. Ullsperger & Debener, 2010). The motivation for using ICA to decompose multi-channel EEG recordings is twofold. On the one hand, ICA can be used to detect and remove artifacts (Jung et al., 2000), on the other hand, it may also be used to unmix brain signals (Makeig et al., 2002). The basic idea of deploying ICA to remove the CFA is to identify components in IC-space, that is, the decomposed signal, that predominantly reflect cardiac activity. To remove the artifact, these ICs can then be omitted when projecting the data back to sensor-space. Identifying ICs reflecting cardiac activity can be done manually, as they usually exhibit a distinctive, asymmetric topography (Ullsperger & Debener, 2010), as well as a time course resembling the

ECG when analyzed in a heartbeat-locked manner. This method has been used successfully (e.g. Al et al., 2020, 2021), sometimes assisted by correlating the IC activation with the ECG (Bury et al., 2019). Several tools have been developed to improve the reproducibility of this, at least to some extent, subjective IC selection process. These tools are either focused on selecting ICs reflecting the CFA specifically (Issa et al., 2019; Tamburro et al., 2019) or offer general assistance in IC classification (Pion-Tonachini et al., 2019; Radüntz et al., 2017; Viola et al., 2009). Another problem related to identifying ICs that reflect the CFA is that ICs are not necessarily process-pure. Particularly in high density EEG recordings, specific sources can be reflected by multiple ICS. Furthermore, the separation of artifacts from brain-related activity is not always satisfactory in such a way that EEG and CFA might still be mixed to some extent in IC-space (Ullsperger & Debener, 2010). These mixed source ICs are more likely when the ratio of degrees of freedom for the decomposition and number of sources of variance is low, for example when the number of EEG channels is low, when the experimental task is complex, or when the recording environment is noisy. This can render the decision about whether an IC represents primarily cardiac or brain activity difficult. The regression-based approach presented here avoids these potential pitfalls by not relying on spatial properties of the signal, estimating the CFA on a single-trial basis and being purely data driven without the subjective selection of components. It is also important to note that the applicability of the proposed method does not rely on a specific timing of experimental stimulation and cardiac phase; the only prerequisite is that the EEG and ECG signals are recorded simultaneously.

There are, however, also downsides in comparison to existing methods. An obvious disadvantage is that the ECG needs to be recorded to perform the regression. Nevertheless, as this approach is most relevant for research investigating the interdependency of cognitive and cardiac processes, the ECG will probably be recorded anyway. A possible limitation is also that we restricted the participants' head movements by using a medical cervical collar, thereby stabilizing the spatial propagation of the CFA. Although many EEG studies aim to stabilize head position, this may hamper the generalizability of our approach. This, however, also applies to ICA-based methods for CFA removal. For future studies, it might be possible to record the head movements by motion tracking systems and use this data as an additional feature in the regression models.

In conclusion, our results show that the proposed method to remove the CFA from the EEG signal using neural network regression is effective. Moreover, it allows for a specific removal of the CFA while preserving

stimulus-induced variance even if temporally related to the activity of the heart. The proposed neural network regression approach can thus provide a replicable and automated method for CFA removal, in particular for studies presenting stimuli time-locked to heartbeats.

AUTHOR CONTRIBUTIONS

Stefan Arnau: Conceptualization; formal analysis; methodology; software; visualization; writing – original draft. **Fariba Sharifian:** Methodology; writing – review and editing. **Edmund Wascher:** Conceptualization; writing – review and editing. **Mauro F Larra:** Conceptualization; methodology; writing – review and editing.

ACKNOWLEDGMENTS

We would like to thank Pia Deltenre and Barbara Foschi for their assistance in collecting the data as well as Tobias Blanke for programming the experiment. Open Access funding enabled and organized by Projekt DEAL.

FUNDING INFORMATION

This research was in part funded by a grant from the German Research Foundation (DFG LA 4830/2-1).

CONFLICT OF INTEREST STATEMENT

The authors declare no conflict of interests.

DATA AVAILABILITY STATEMENT

The code used for preprocessing, cleaning, and analyzing the data presented in this study can be found at: https://github.com/stefanarnau/cfa_removal. This Python code can be used to perform the neural-network-based approach for removing the CFA from the EEG on a channel-by-channel basis. The EEG data used in this study can be obtained from the public OSF repository at <https://osf.io/8b6mt/>.

ORCID

Stefan Arnau  <https://orcid.org/0000-0002-8858-4287>

Edmund Wascher  <https://orcid.org/0000-0003-3616-9767>

Mauro F. Larra  <https://orcid.org/0000-0001-5721-0920>

REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D. G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., ... Zheng, X. (2015). *TensorFlow: A system for large-scale machine learning*. 21.
- Al, E., Iliopoulos, F., Forschack, N., Nierhaus, T., Grund, M., Motyka, P., Gaebler, M., Nikulin, V. V., & Villringer, A. (2020). Heart-brain interactions shape somatosensory perception and evoked potentials. *Proceedings of the National Academy of Sciences*, 117(19), 10575–10584. <https://doi.org/10.1073/pnas.1915629117>

- Al, E., Iliopoulos, F., Nikulin, V. V., & Villringer, A. (2021). Heartbeat and somatosensory perception. *NeuroImage*, 238, 118247. <https://doi.org/10.1016/j.neuroimage.2021.118247>
- Azevedo, R. T., Garfinkel, S. N., Critchley, H. D., & Tsakiris, M. (2017). Cardioafferent activity modulates the expression of racial stereotypes. *Nature Communications*, 8(1), 13854. <https://doi.org/10.1038/ncomms13854>
- Azzalini, D., Rebollo, I., & Tallon-Baudry, C. (2019). Visceral signals shape brain dynamics and cognition. *Trends in Cognitive Sciences*, 23(6), 488–509. <https://doi.org/10.1016/j.tics.2019.03.007>
- Bury, G., García-Huésca, M., Bhattacharya, J., & Ruiz, M. H. (2019). Cardioafferent activity modulates early neural signature of error detection during skilled performance. *NeuroImage*, 199, 704–717. <https://doi.org/10.1016/j.neuroimage.2019.04.043>
- Chollet, F., et al. (2015). Keras.GitHub. <https://github.com/fchollet/keras>
- Clevert, D.-A., Unterthiner, T., & Hochreiter, S. (2016). Fast and accurate deep network learning by exponential linear units (ELUs). *ArXiv:1511.07289 [Cs]*. <http://arxiv.org/abs/1511.07289>
- Cohen, J. (1992). A power primer. *Psychological Bulletin*, 112, 155–159. <https://doi.org/10.1037/0033-2909.112.1.155>
- Critchley, H. D., & Harrison, N. A. (2013). Visceral influences on brain and behavior. *Neuron*, 77(4), 624–638. <https://doi.org/10.1016/j.neuron.2013.02.008>
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134(1), 9–21. <https://doi.org/10.1016/j.jneumeth.2003.10.009>
- Dirlich, G., Vogl, L., Plaschke, M., & Strian, F. (1997). Cardiac field effects on the EEG. *Electroencephalography and Clinical Neurophysiology*, 102(4), 307–315. [https://doi.org/10.1016/S0013-4694\(96\)96506-2](https://doi.org/10.1016/S0013-4694(96)96506-2)
- Dozat, T. (2016). *Incorporating nesterov momentum into ADAM*. 4.
- Garfinkel, S. N., & Critchley, H. D. (2016). Threat and the body: How the heart supports fear processing. *Trends in Cognitive Sciences*, 20(1), 34–46. <https://doi.org/10.1016/j.tics.2015.10.005>
- Garfinkel, S. N., Minati, L., Gray, M. A., Seth, A. K., Dolan, R. J., & Critchley, H. D. (2014). Fear from the heart: Sensitivity to fear stimuli depends on individual heartbeats. *Journal of Neuroscience*, 34(19), 6573–6582. <https://doi.org/10.1523/JNEUROSCI.3507-13.2014>
- Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., & Hämäläinen, M. (2013). *MEG and EEG data analysis with MNE-Python*. 7 <https://doi.org/10.3389/fnins.2013.00267>
- Gray, M. A., Minati, L., Paoletti, G., & Critchley, H. D. (2010). Baroreceptor activation attenuates attentional effects on pain-evoked potentials. *Pain*, 151(3), 853–861. <https://doi.org/10.1016/j.pain.2010.09.028>
- He, K., Zhang, X., Ren, S., & Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. *ArXiv:1502.01852 [Cs]*. <http://arxiv.org/abs/1502.01852>
- Hjorth, B. (1975). An on-line transformation of EEG scalp potentials into orthogonal source derivations. *Electroencephalography and Clinical Neurophysiology*, 39(5), 526–530. [https://doi.org/10.1016/0013-4694\(75\)90056-5](https://doi.org/10.1016/0013-4694(75)90056-5)
- Hyvärinen, A., Karhunen, J., & Oja, E. (2001). *Independent component analysis*. John Wiley & Sons, Inc. <https://doi.org/10.1002/047121317>
- Issa, M. F., Tuboly, G., Kozmann, G., & Juhasz, Z. (2019). Automatic ECG artefact removal from EEG signals. *Measurement Science Review*, 19(3), 101–108. <https://doi.org/10.2478/msr-2019-0016>
- Jung, T.-P., Makeig, S., Humphries, C., Lee, T.-W., McKeown, M. J., Iragui, V., & Sejnowski, T. J. (2000). Removing electroencephalographic artifacts by blind source separation. *Psychophysiology*, 37(2), 163–178. <https://doi.org/10.1111/1469-8986.3720163>
- Kayser, J., & Tenke, C. E. (2006). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clinical Neurophysiology*, 117(2), 348–368. <https://doi.org/10.1016/j.clinph.2005.08.034>
- Larra, M. F., Finke, J. B., Wascher, E., & Schächinger, H. (2020). Disentangling sensorimotor and cognitive cardioafferent effects: A cardiac-cycle-time study on spatial stimulus-response compatibility. *Scientific Reports*, 10(1), 4059. <https://doi.org/10.1038/s41598-020-61068-1>
- Lopes, F., Leal, A., Medeiros, J., Pinto, M. F., Dourado, A., Dumpelmann, M., & Teixeira, C. (2021). Automatic electroencephalogram artifact removal using deep convolutional neural networks. *IEEE Access*, 9, 149955–149970. <https://doi.org/10.1109/ACCESS.2021.3125728>
- Luft, C. D. B., & Bhattacharya, J. (2015). Aroused with heart: Modulation of heartbeat evoked potential by arousal induction and its oscillatory correlates. *Scientific Reports*, 5(1), 15717. <https://doi.org/10.1038/srep15717>
- Makeig, S., Westerfield, M., Jung, T.-P., Enghoff, S., Townsend, J., Courchesne, E., & Sejnowski, T. J. (2002). Dynamic brain sources of visual evoked responses. *Science*, 295(5555), 690–694. <https://doi.org/10.1126/science.1066168>
- Merlin Praveena, D., Angelin Sarah, D., & Thomas George, S. (2022). Deep learning techniques for EEG signal applications—A review. *IETE Journal of Research*, 68(4), 3030–3037. <https://doi.org/10.1080/03772063.2020.1749143>
- Montoya, P., Schandry, R., & Müller, A. (1993). Heartbeat evoked potentials (HEP): Topography and influence of cardiac awareness and focus of attention. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, 88(3), 163–172. [https://doi.org/10.1016/0168-5597\(93\)90001-6](https://doi.org/10.1016/0168-5597(93)90001-6)
- Mordkoff, J. T. (2019). A simple method for removing bias from a popular measure of standardized effect size: Adjusted partial eta squared. *Advances in Methods and Practices in Psychological Science*, 2(3), 228–232. <https://doi.org/10.1177/2515245919855053>
- Pion-Tonachini, L., Kreutz-Delgado, K., & Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, 198, 181–197. <https://doi.org/10.1016/j.neuroimage.2019.05.026>
- Pollatos, O., & Schandry, R. (2004). Accuracy of heartbeat perception is reflected in the amplitude of the heartbeat-evoked brain potential: Heartbeat-evoked potential and heartbeat perception. *Psychophysiology*, 41(3), 476–482. <https://doi.org/10.1111/1469-8986.2004.00170.x>
- Pramme, L., Larra, M. F., Schächinger, H., & Frings, C. (2016). Cardiac cycle time effects on selection efficiency in vision: Baroreceptor activity and visual attention. *Psychophysiology*, 53(11), 1702–1711. <https://doi.org/10.1111/psyp.12728>
- Radüntz, T., Scouten, J., Hochmuth, O., & Meffert, B. (2017). Automated EEG artifact elimination by applying machine

- learning algorithms to ICA-based features. *Journal of Neural Engineering*, 14(4), 046004. <https://doi.org/10.1088/1741-2552/aa69d1>
- Schulz, A., Lass-Hennemann, J., Nees, F., Blumenthal, T. D., Berger, W., & Schachinger, H. (2009). Cardiac modulation of startle eye blink. *Psychophysiology*, 46(2), 234–240. <https://doi.org/10.1111/j.1469-8986.2008.00768.x>
- Tamburro, G., Stone, D. B., & Comani, S. (2019). Automatic removal of cardiac interference (ARCI): A new approach for EEG data. *Frontiers in Neuroscience*, 13, 441. <https://doi.org/10.3389/fnins.2019.00441>
- Terhaar, J., Viola, F. C., Bär, K.-J., & Debener, S. (2012). Heartbeat evoked potentials mirror altered body perception in depressed patients. *Clinical Neurophysiology*, 123(10), 1950–1957. <https://doi.org/10.1016/j.clinph.2012.02.086>
- Ullsperger, M., & Debener, S. (Eds.). (2010). *Simultaneous EEG and fMRI: Recording, analysis, and application*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780195372731.001.0001>
- van Elk, M., Lenggenhager, B., Heydrich, L., & Blanke, O. (2014). Suppression of the auditory N1-component for heartbeat-related sounds reflects interoceptive predictive coding. *Biological Psychology*, 99, 172–182. <https://doi.org/10.1016/j.biopsycho.2014.03.004>
- Viola, F., Thorne, J., Edmonds, B., Schneider, T., Eichele, T., & Debener, S. (2009). Semi-automatic identification of independent components representing EEG artifact. *Clinical Neurophysiology*, 120(5), 868–877. <https://doi.org/10.1016/j.clinph.2009.01.015>

How to cite this article: Arnau, S., Sharifian, F., Wascher, E., & Larra, M. F. (2023). Removing the cardiac field artifact from the EEG using neural network regression. *Psychophysiology*, 00, e14323. <https://doi.org/10.1111/psyp.14323>