Making Noise - Improving Seismocardiography Based Heart Analysis With Denoising Autoencoders

Jonas Burian*
Helmut Tödtmann*
jonas.burian@aalto.fi
helmut.toedtmann@igd-r.fraunhofer.de
Fraunhofer IGD
Rostock, Germany

Marian Haescher
Fraunhofer IGD
Rostock, Germany
marian.haescher@igd-r.fraunhofer.de

Mario Aehnelt
Fraunhofer IGD
Rostock, Germany
mario.aehnelt@igd-r.fraunhofer.de

Arjan Kuijper
Fraunhofer IGD
Darmstadt, Germany
arjan.kuijper@igd.fraunhofer.de

ABSTRACT

Seismocardiography is a method commonly used to monitor and prevent cardiovascular diseases. However, noise and artifacts in the signals often interfere with the assessment of cardiac health and the analysis of the signal morphology. Therefore, this work presents a new approach to denoise seismocardiography signals by applying fully convolutional denoising autoencoders. In order to investigate the suitability and robustness of this approach, a series of experiments have been carried out with respect to the optimal configuration for the denoising task and a comparison with wavelet denoising as a traditional approach. Furthermore, the practical applicability of the method is tested with the use case of transforming noisy seismocardiography signals into electrocardiography signals. Our approach using autoencoders outperforms the commonly used wavelet denoising. Additionally, we demonstrate that the widespread usage of Butterworth filters may not only be unnecessary but even detrimental. Finally, the generalizability of the method is verified on unseen data. With those combined improvements in noise reduction, the assessment of cardiac health using seismocardiography in the presence of noise may be facilitated in the future.

CCS CONCEPTS

• **Applied computing** → *Health informatics*; *Medical records systems*; *Healthcare administration*; *Telemedicine*; • **Computing methodologies** → **Machine learning**; • **Biomedical informatics and computational biology**;

KEYWORDS

seismocardiography, cardiac health assessment, signal denoising, convolutional autoencoder, signal processing

ACM Reference Format:
Jonas Burian, Helmut Tödtmann, Marian Haescher, Mario Aehnelt, and Arjan Kuijper. 2023. Making Noise - Improving Seismocardiography Based Heart Analysis With Denoising Autoencoders. In *8th international Workshop on Sensor-Based Activity Recognition and Artificial Intelligence (iWOAR 2023), September 21–22, 2023, Lübeck, Germany*. ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3615834.3615847

*Both authors contributed equally to this research.

iWOAR 2023, September 21–22, 2023, Lübeck, Germany

ACM ISBN 979-8-4007-0816-9/23/09…$15.00
https://doi.org/10.1145/3615834.3615847

1 INTRODUCTION

Cardiovascular Diseases (CVDs) are the number one cause of mortality worldwide. According to the World Health Organization (WHO), 17.9 million people died from CVDs in 2019, representing more than half of deaths from noncommunicable diseases, which account for 74 % of all deaths [15]. This emphasizes the importance of monitoring the cardiac health of people, as early detection of such diseases allows for an early treatment [14, 15]. The monitoring can be done, for example, by using a technique called Seismocardiography (SCG), which allows the non-invasive measurement of the vibrations induced by the heart at a localized region on the chest [18, 23].

However, SCG signals are often contaminated by noise and artifacts resulting from the mechano-electronics of the sensor used, the subject's motion, and vibrations from the environment [17, 19]. This in turn implies the need to denoise SCG signals in order to extract the precise signal morphology required to assess the cardiac health. In the past, more traditional techniques such as bandpass filters, Wavelet Denoising (WD) or Empirical Mode Decomposition (EMD) have been used to denoise SCG signals. Over the past decades, however, deep learning methods, and in particular Autoencoders (AEs), have achieved very good results in denoising signals such as Electrocardiography (ECG) signals [3], another popular method used to monitor cardiac health.

In this work, we present a new method for denoising SCG signals using Denoising Autoencoders (DAEs) based on Fully Convolutional Neural Networks (FCNs), a structure that has shown promise in the past for denoising other signals [3]. In addition, the practical

applicability of the approach is evaluated on the use case of transforming noisy SCG to ECG signals, as proposed by [8]. In order to investigate the suitability of this method and to find an optimal configuration for the denoising task, four different experiments were performed:

(1) Investigating the effects of applying Butterworth bandpass filters at different points in the denoising process when using FCN-based DAEs
(2) Investigating different distributions of noise levels for training the AEs
(3) Comparing AE-based denoising with WD as a traditional approach
(4) Evaluating the real-world applicability of the approach with the use case of transforming noisy SCG to ECG signals

The rest of this paper is structured as follows: First, previous work in the field of signal denoising with SCG signals is presented, followed by the methodology of the four experiments, which includes the datasets used, the pre-processing of the data and the modeling of noise, the AEs, as well as the experimental setup and the evaluation metrics. Finally, the results of the experiments are presented and discussed.

2 RELATED WORK

Most commonly used to remove noise and artifacts from SCG signals are bandpass, low-, highpass, and notch filters [17, 19]. They are used primarily to remove artifacts caused by body motion and respiration. Typically bandpass filters are infinite response filters with the butterworth filter as most prominent example, using cutoff frequencies of 0.8–35 Hz, 0.8-25 Hz and 0.5–40 Hz, respectively. For example, a high-pass filter can remove baseline wandering caused by respiration [5], while bandpass filters can reduce speech artifacts and the low-frequency components of walking-induced noise [11]. Such analog filters can also be combined with ensemble averaging for a possible additional denoising effect [9] – or the latter method is used on its own [5]. Savitzky Golay based polynomial smoothing has been applied to improve heart sound extraction from seismocardiograms [16]. Indirect denoising via combining the outputs of multiple connected accelerometers has also proven to increase the overall signal quality [13]. For SCG related ballistocardiograms, ICA showed promising results for denoising [2], the authors enhance biomedical signal filtering by proposing a threefold independent component analysis-based approach, yielding improved performance across various test scenarios. For walking-induced noise in particular, EMD and Ensemble Empirical Mode Decomposition (EEMD) have been widely used in the past [9, 19], also in combination with additional corruption removal [12]. Furthermore, normalized least mean square adaptive filters were applied to reduce these types of artifacts as well [22], increasing heartbeat detection rate from 91% to 98%. Finally, in a comparison between wavelet transforms, EMD, and adaptive and morphological filtering, wavelet thresholding was found to be the most suitable for denoising SCG signals corrupted by White Gaussian Noise (WGN) [1], it scored the best values in all relevant metrics (MSE, MAE, PSNR, xcorr) and was comparable in terms of best run time.

3 METHODOLOGY

In this section, the underlying methodology of the experiments is presented. This includes not only a description of the experiments and the evaluation metrics themselves, but also a presentation of the datasets used and their pre-processing, the noise model applied, and the architecture of the AEs used.

3.1 Datasets

In total, the combination of three different datasets is used in the experiments for training the AEs and for the evaluation. The first dataset is the Combined Measurement of ECG, Breathing and Seismocardiogram (CEBS, here: D1) [6, 7]. Included are recordings from 20 "presumably healthy" subjects (7 females and 13 males) with an average age of 24.4 ± 3.1 years. The SCG signals were recorded using a triaxial accelerometer (LIS344ALH, ST Microelectronics) on channel 4 of a Biopac MP39 data acquisition system (Santa Barbara, CA, USA), while channels 1 and 2 were used to measure the two-lead ECG signal. The recordings were made with the subjects lying very still in a supine position to ensure a calm environment.

The second dataset used (D2) also contains ECG, SCG and Gyrocardiography (GCG) recordings from 100 different people (41 female and 59 male) with heart valve diseases and an average age of 68 ± 14 years [21]. To record the data, a Shimmer 3 ECG module (Shimmer Sensing, United Kingdom) was used to record 3 ECG channels but also the SCG using a triaxial MEMS accelerometer. For the experiments, however, only the z-axis, i.e. the axis in a dorsoventral direction, was used. The patients lay in a supine position, were asked to stay awake and to breathe normally.

Like the second dataset, the third dataset (D3) contains ECG, SCG, and GCG measurements, where the SCG signals were recorded using a triaxial capacitive digital accelerometer (Freescale Semiconductor, MMA8451Q, Austin, TX, USA), but again only the z-axis was used in the experiments [10]. The 29 male subjects were aged 29 ± 5 years on average and lay in a supine position on their left or right side.

3.2 Pre-processing and Noise Model

To use the data for training and evaluation, it was resampled to 100 Hz, corresponding to the acceleration sampling capabilities of mobile devices between 100 and 200 Hz [8]. In a next step, a fourth-order Butterworth bandpass filter with cutoff frequencies of 5 Hz and 30 Hz was applied to the SCG samples to remove noise such as baseline wandering present in the data [6, 8]. To remove unusual peaks and outliers, the z-score of each data point was calculated and any data point with a value above a threshold of 7 was set to the mean of the signal. In a next step, a moving window with an offset of 64 was applied to divide the SCG and ECG data into samples of 512 data points. The SCG data was then normalized between -1 and 1.

In the experiments, additive WGN was used as the noise model to train the DAEs and to test the denoising performance of the different denoising methods. The noise was added to the clean SCG samples with a specific signal-to-noise ratio (SNR) according to a given noise distribution, i.e., a set of SNRs from which one value is randomly selected for each sample. These noise distributions are -2.5, 0, 2.5 dB (WGN1); -2.5, 0, 2.5, 2.5, 5, 7.5 dB (WGN2); -2.5, 0, 2.5,

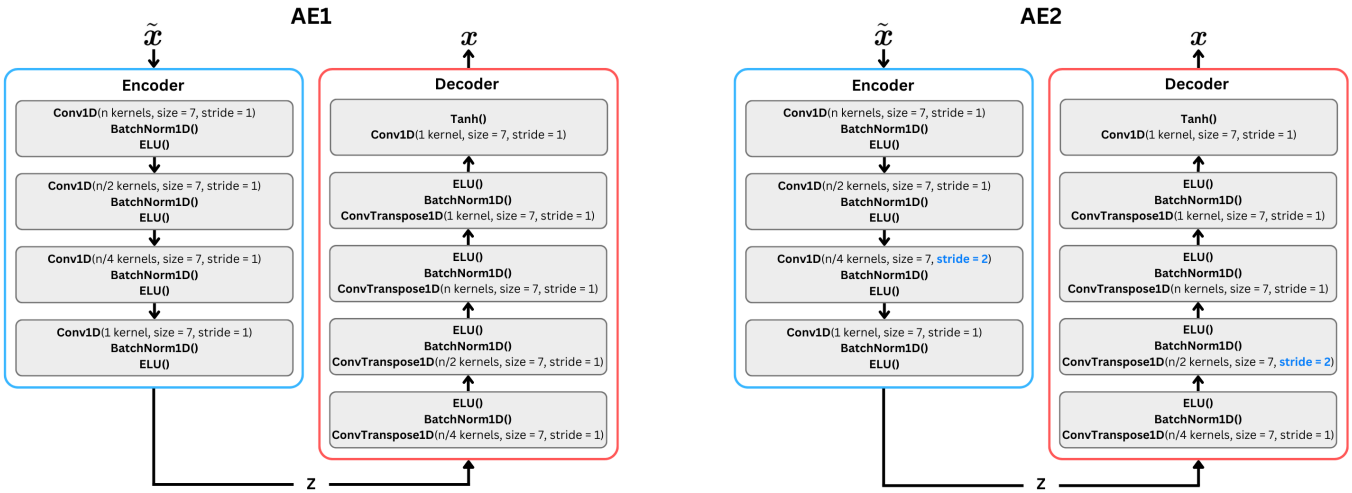

Figure 1: Architectures of AE1 (left) and AE2 (right).

7.5, 10, 12.5 dB (WGN3) and -2.5, 0, 2.5, 17.5, 20, 22.5 dB (WGN4) for training and -1, 0, 2, 5, 7 dB (WGN5) for testing, following and building on the example in [3]. The different training distributions will be evaluated in one of the later experiments.

3.3 Autoencoder Architectures

Autoencoder 1 (AE1) in the experiments, as shown in Figure 1, consists of four one-dimensional convolutional layers in the encoder and four transposed convolutional layers in the decoder, each having a kernel size of 7 and a stride of 1. The first layer of the encoder has n kernels, with this number being halved in each subsequent layer and the last layer having only one kernel. This number of kernels will later be set to 128, which was determined by hyperparameter tuning. Each layer is followed by a one-dimensional batch normalization layer and the Exponential Linear Unit (*ELU*) activation function with the parameter α, which defines the saturation of the function for negative inputs [4], set to its default value of 1. The *ELU* is defined as

$$g(z) = z \; if \; z > 0$$
$$\alpha(exp(z) - 1) \; if \; z <= 0$$

and especially when being compared to Rectified Linear Units (ReLUs), allows for a faster learning as negative input values are possible and therefore the activation means are pushed closer to zero [4]. *ELUs* also decrease the information and variation propagated forward because of a small derivative which comes from a saturation to negative values for smaller inputs [4]. For the decoder, which is build inversely-symmetrical to the encoder, the number of kernels is doubled in each of the first three convolutional layers from $n/4$ to n. The last transposed convolutional layer has only 1 kernel. Again, a one dimensional batch normalization is used after each layer followed by the *ELU* activation function. The output layer is in the form of a one-dimensional convolutional layer with 1 kernel with a size of 7 and a stride of 1. This is followed by the *Tanh* activation function.

The structure of Autoencoder 2 (AE2), as also shown in Figure 1, is almost identical to AE1, except that the penultimate layer of the encoder and the second layer of the decoder have a stride of 2 instead of 1. This results in the input size being almost halved in the code layer. The number of kernels, n, for AE2 is set to 128 for the denoising task and to 256 for the transformation task from SCG to ECG signals, following the suggestion in [8].

3.4 Experiments and Evaluation Metrics

Four different experiments are conducted: In the first experiment, the use of a fourth-order Butterworth bandpass filter with cutoff frequencies of 5 and 30 Hz [6, 8] is investigated for its use in combination with AE2 as a denoising method, i.e. the filter is applied on the training and test data for AE2 only, on the denoised results by AE2 only, and on both the training/test data and the results. This is then compared to no use of bandpass filters, i.e. the the sole use of AE2 to denoise the signals. The training data is corrupted with WGN3, while the test data is corrupted with WGN5.

The second experiment examines different noise distributions for corrupting the training data of AE2. These are the previously presented WGN1 to WGN4, while training with uncorrupted data is also tested. The test data is always corrupted with the test distribution WGN5. In this experiment, no bandpass filter is used on the training/test data or on the results.

The third experiment serves the purpose of comparing the AE-based denoising methods together with WD as a traditional approach that was previously found to be well suited for denoising SCG signals [1]. AE1 and AE2 are trained on WGN3-corrupted data, while testing for both the AEs and WD is performed on data corrupted with WGN5. No bandpass filters are used on the training/test data or on the results. For WD, soft thresholding is used together with Bayesian wavelet shrinkage. In addition, the robustness and generalizability of the AE that performs best is verified by training it on the full D1 dataset with the data of the first and second subject omitted and testing it on the data of these two omitted subjects only. The performance is then compared to training the AE on the

full D1 dataset, i.e. including subjects one and two, and testing it again on the data from these two subjects only. Only the D1 dataset is used because the AE should be prevented from learning distinct features from different datasets. WGN3 and WGN5 are used again to corrupt the training and test data, respectively.

The last experiment investigates the practical implications of the best denoising method from the third experiment on the use case of transforming noisy SCG signals into ECG signals, as proposed by [8] for clean SCG signals. That is, a comparison is made between transforming noisy SCG signals to ECG signals and transforming denoised SCG signals to ECG signals using the best denoising method from the third experiment. For the transformation model, AE2 is trained on clean SCG signals using the ECG signals as ground truth. Since it was found that the transformation performance deteriorates significantly when all datasets are used, only the D1 dataset is used for this experiment.

The denoising performance of the different denoising methods is evaluated using three evaluation metrics: First, the SNR improvement (SNR_{imp}) is used as the difference of the SNR after and before denoising ($SNR_{imp} = SNR_{after} - SNR_{before}$), as proposed by [3], to measure the overall denoising performance. The SNR can be defined as

$$SNR = 10\ log_{10} \frac{\sum_{n=1}^{N} x_i^2}{\sum_{n=1}^{N} (x_i - \hat{x_i})^2} \tag{1}$$

where N denotes the number of datapoints in the signal, x_i a clean signal datapoint and $\hat{x_i}$ the corresponding noisy datapoint or predicted or denoised output.

A second evaluation criterion used is the Root Mean Squared Error (RMSE) between the denoised and clean signal to measure the accuracy of the denoising, defined as $RMSE = \sqrt{MSE}$, with the Mean Squared Error (MSE) being

$$MSE = \frac{1}{N} \sum_{n=1}^{N} (x_i - \hat{x_i})^2 \tag{2}$$

Lastly, the Pearson correlation coefficient (pcorr), defined as

$$r = \frac{\sum_{i=1}^{N} \hat{x}_i x_i - N\bar{\hat{x}}\bar{x}}{\sqrt{(\sum_{i=1}^{N} \hat{x}_i^2 - N\bar{\hat{x}}^2)(\sum_{i=1}^{N} x_i^2 - N\bar{x}^2)}} \tag{3}$$

with $\bar{z} = \frac{1}{N} \sum_{i=1}^{N} z_i$, is used to measure the preservation in structural integrity between the denoised and clean signal. These three metrics are applied to experiments one through three. However, for the last experiment, only the RMSE and the pcorr are used, since the focus of this experiment is not to evaluate the denoising performance, but the transformation capabilities from SCG to ECG.

4 RESULTS

The results of the first experiment, as also displayed in Figure 2, show that the best performance, i.e. the highest overall SNR, the lowest RMSE and the highest pcorr, is obtained when no bandpass filter is used in the denoising process with AE2 trained on WGN3, while the overall worst performance is given when the bandpass filter is applied to the train/test data or to the train/test data and the results. However, with regards to the pcorr, slightly better results of 1.16 % were obtained when a filter was applied to the results compared to no application to the results. A significant feature of the results is also that for all four cases, the results get worse for the SNR improvement, but better for the RMSE and the pcorr with less noise corruption of the signals. Finally, the Standard Deviation (SD) for all three evaluation metrics is almost the same for the four evaluation cases.

In the second experiment, the best results are achieved by AE2 trained on the data corrupted with WGN3, while the worst results are obtained with the first noise distribution, as shown in Figure 3. It can be seen that the results get better from the first to the third distribution, whereas AE2 trained on WGN4 achieves slightly worse results compared to WGN3. Furthermore, the results show that the AE trained on clean data achieves better results than the AE trained on WGN1. Regarding the SD, the clean training data and the data corrupted by WGN3 lead to significantly lower values for the SNR improvement, while for the RMSE, the clean training data and for the pcorr, WGN1 gives the worst results.

For the results of the third experiment, as shown in Figure 4 for the comparison part and in Figure 5 for the verification part of the experiment, it is significant that the best results were obtained by AE2, followed by WD and then AE1. The two AEs were trained with WGN3 corrupted training data and without the use of bandpass filters, based on the results of the previous experiments. Results for AE2 are especially better for more corrupted signals, but almost equal to WD for less corrupted signals. Specifically, AE2 achieved 50 % better results than WD for the SNR improvement, while the results for the RMSE and the pcorr were 25 % and 7.5 % better, respectively. Looking at the SD, AE2 and WD both show significantly lower results for SNR improvement, while AE1 gives the best results for RMSE and AE2 for the pcorr. In addition, the average CPU processing time for the 5,120 ms long sequences was 10.50 ± 5.77 ms and 10.23 ± 5.97 ms for AE1 and AE2, respectively, while the processing time for the signals with WD was only 0.62 ± 1.35 ms on average. Finally, the verification part of the experiment shows slightly better results for AE2 trained on the full D1 dataset, but lower results for the SD in terms of the SNR improvement and pcorr for training on the dataset with omitted subjects.

Training AE2 on the D1 dataset to transform clean SCG signals into ECG signals in the last experiment resulted in a pcorr of 0.93 ± 0.097 with a RMSE of 0.065 ± 0.03 when the ECG signals were normalized. However, without normalizing the ECG signals, a higher pcorr of 0.98 ± 0.027 and a lower RMSE of 0.022 ± 0.015 could be achieved. Therefore, the ECG signals will not be normalized in the remaining experiment. Regarding the transformation from noisy SCG signals to ECG signals, considerably better results could be achieved when AE2 was used as a denoising method before the transformation to ECG signals. For a noise level of SNR -1 dB, the pcorr increased more than 10 times while the RMSE decreased by 27.27 %. At a noise level of SNR 7 dB, the changes were not as pronounced, but still showed significant improvements.

5 DISCUSSION

The results of the first experiment showed that bandpass filters should not be used in the denoising process with AE2, as their application slightly degrades the overall denoising performance. Although the application of bandpass filters on the denoised signals led to slightly better results for the pcorr, they should still not

be used, mainly because of the higher computational time, which may be a hindrance in applications where real-time processing or low computational time is required. The overall inferior denoising performance may be due to two main reasons: first, as certain frequencies are suppressed, the use of bandpass filters may destroy relationships between the SCG and noise signal morphology, so that the AE cannot learn these relationships well enough. This also means that the use of such filters hinders the data augmentation process. In turn, the generalizability and robustness of the AE may be compromised because of less (diverse) training data.

For the second experiment, the results suggest that AE2 should be trained with a noise distribution containing both very noisy and less noisy signals. On the one hand, more noisy signals help the AE to learn the discriminative features of the noise and the signal itself [20]. On the other hand, the less noisy signals contribute significantly to the denoising performance of less but also more noisy signals. In fact, it seems that the AE is also able to denoise inherently, so that less noisy train samples contribute not only to the denoising performance of less noisy test samples, but also to the denoising performance of more noisy test signals, as can be

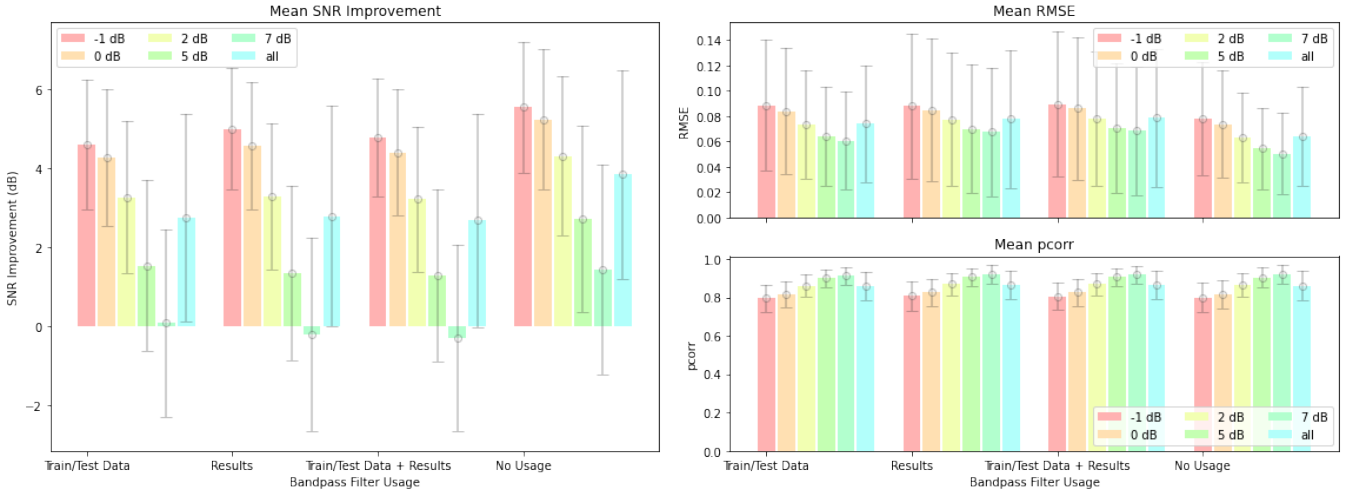

Figure 2: Results of the first experiment. Shown are the mean SNR improvement, RMSE and pcorr with their respective SD for denoising SCG signals corrupted with WGN5. The signals were denoised using AE2 trained on WGN3 and a bandpass filter was applied to the train/test data only, to the results only, to the train/test data and the results, and also no bandpass filter was used. Each SNR of WGN5 is shown in dB, while 'all' indicates the results over all SNRs of WGN5. See table 1.

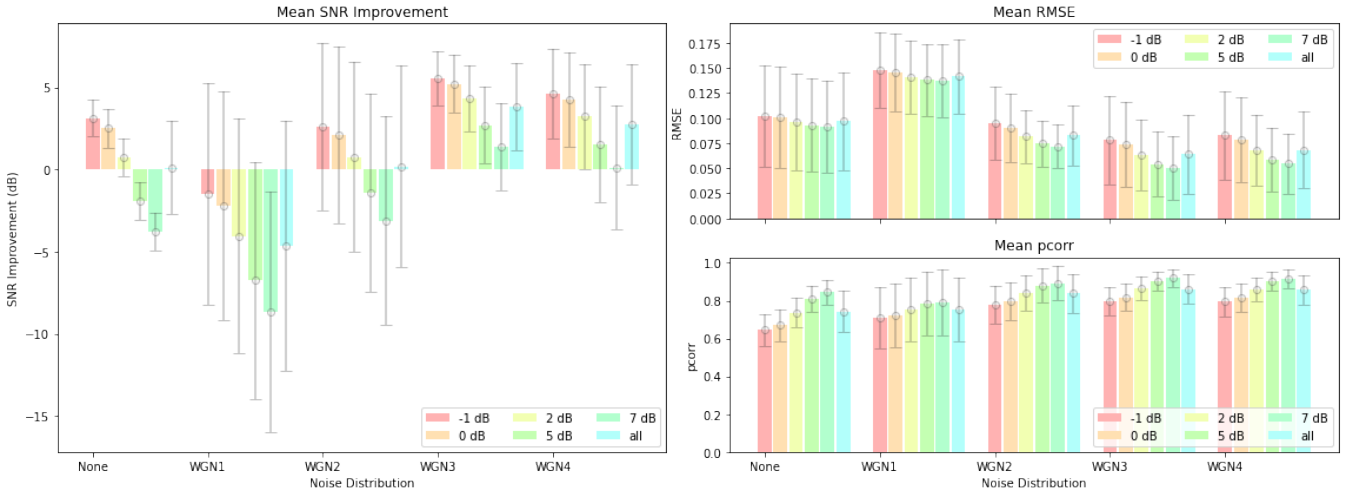

Figure 3: Results of the second experiment. Shown are the mean SNR improvement, RMSE and pcorr with their respective SD for denoising SCG signals corrupted with WGN5. The signals have been denoised by AE2 trained on a dataset corrupted with WGN1 to WGN4, while the denoising by AE2 trained on uncorrupted data (None) is also shown. Each SNR of WGN5 is shown in dB, while 'all' indicates the results over all SNRs of WGN5. See table 4.

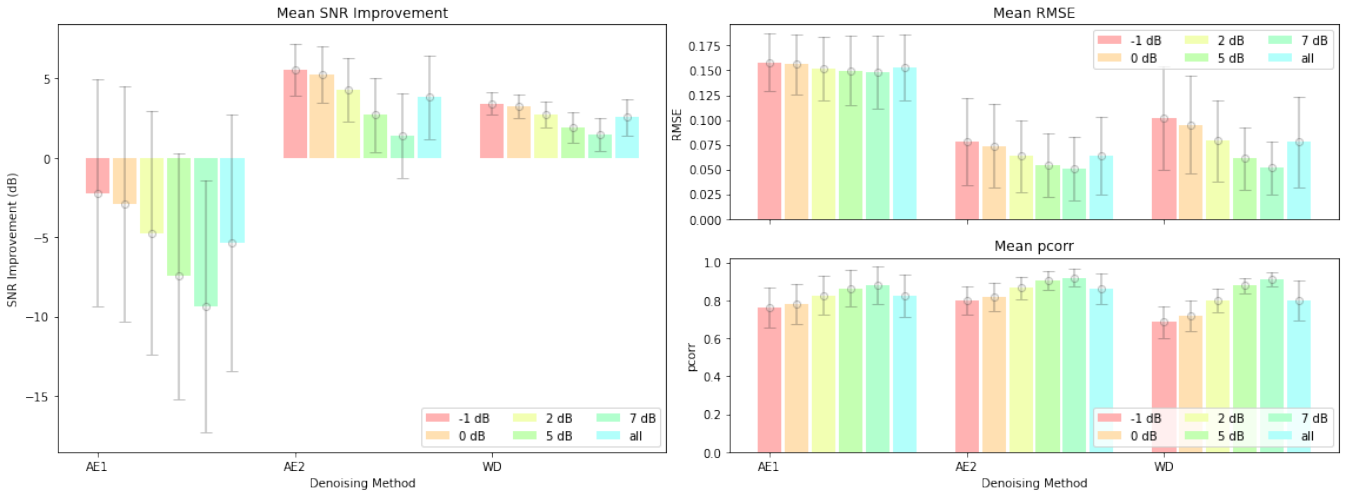

Figure 4: Results of the first part of the third experiment. Shown are the mean SNR improvement, RMSE and pcorr with their respective SD for denoising SCG signals corrupted with WGN5. The signals were then denoised using AE1 and AE2 trained on a dataset corrupted with WGN3, but also with WD as a traditional denoising approach. Each SNR of WGN5 is shown in dB, while 'all' indicates the results over all SNRs of WGN5. See table 7.

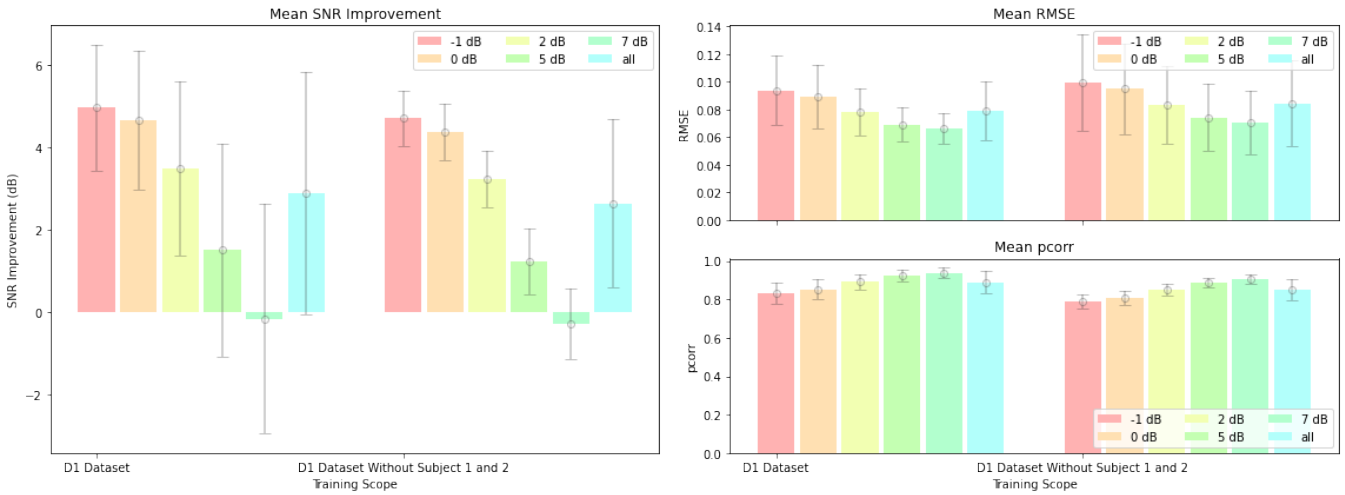

Figure 5: Results of the verification part of the third experiment. Shown are the mean SNR improvement, RMSE and pcorr with their respective SD for denoising the SCG signals of subjects 1 and 2 from the D1 dataset corrupted with WGN5. The results compare denoising with AE2 trained on the full D1 dataset and on the D1 dataset with subjects 1 and 2 omitted, where in both cases the signals were corrupted with WGN3. Each SNR of WGN5 is shown in dB, while 'all' indicates the results over all SNRs of WGN5. See table 10.

clearly observed in Figure 3. Finally, it should be noted that the inclusion of less noisy signals in the training data (WGN3/4) results in a lower SD.

In the third experiment, AE2 proved to provide the best overall results in denoising SCG signals when compared to denoising with WD and AE1. This is particularly evident when AE2 is presented with more noisy signals, while it performs comparably to WD when denoising less noisy signals. The worse results for AE1 can be attributed to the lower encoding rate present, as the rest of the structure of the two AEs is identical. In addition, AE2 shows almost equal performance on subjects it has seen before during training and on subjects it has not seen during training, indicating that it has good generalization capabilities by learning the properties of different SCG signal morphologies and the properties of the added noise, rather than just 'memorizing' the training data. This makes AE2 suitable for real-world applications where, in most cases, the subjects in the training data are not examined.

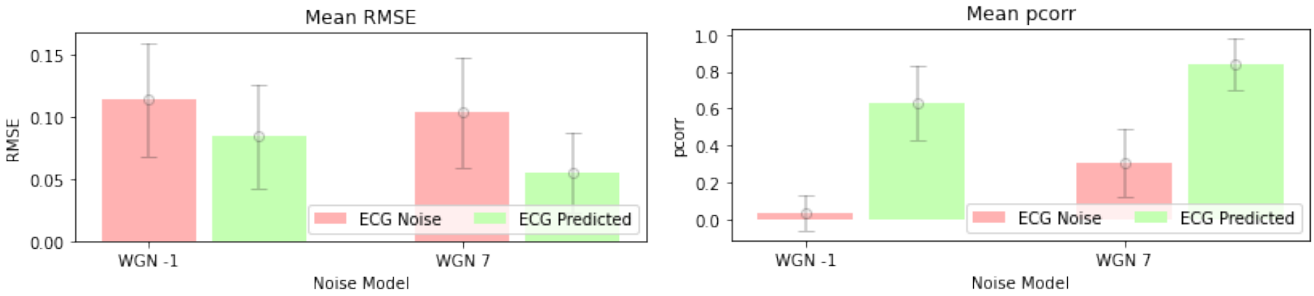

Figure 6: Results of the fourth experiment. Shown are the mean RMSE and pcorr with their respective SD for the transformation of noisy SCG signals from the D1 dataset corrupted with an SNR of -1 dB and 7 dB. AE2 was trained for the transformation task on the D1 dataset with clean SCG/ECG samples. A comparison is made between the transformation from the noisy samples (ECG Noise) and the transformation from denoised signals using AE2 trained on all three datasets corrupted with WGN3 as the denoising method (ECG Predicted).

The last experiment shows better results for the transformation of clean SCG to ECG signals when no normalization is applied to the ECG signals, possibly indicating that the normalization destroys important relationships about the scaling between the SCG and ECG signals and their respective signal morphologies. Regarding the transformation from noisy SCG signals, no useful ECG signal can be extracted, as indicated by the low pcorr and high RMSE. However, after denoising the signals with AE2, the results were substantially improved, demonstrating the practical applicability of using AE2 as a denoising approach for the SCG/ECG signal transformation use case. Nevertheless, the clinical relevance and morphological usability need to be assessed by domain experts in the future, as also done in [8], through a morphological and rhythmic evaluation.

With regard to a real-time application, it can be said that this is easily possible due to the calculation times given in the previous section; the bottleneck in this case is rather the connection and transfer times from mobile devices to the service in the cloud, which exceed the calculation times by several orders of magnitude.

6 CONCLUSION AND FUTURE WORK

In this paper, we presented a new approach to denoise SCG signals using FCN-based DAE. Our approach outperforms WD as another well-suited approach that has been found to deliver favorable results in the past [1]. In particular, the AE with the higher encoding rate showed the best results and was able to denoise signals from patients not included in the training data with almost the same denoising performance. Other findings include that the AE should be trained with less and more noisy signals, and that commonly used bandpass filters can hinder the denoising performance of the AE. Finally, the practical applicability of this approach was demonstrated with the use case of transforming noisy SCG to ECG signals by significantly increasing the transformation performance. These results may help in the future to apply SCG signals to a wider range of applications and environments where noise is present, such as when the subjects move around during the recording.

In future work, more elaborate datasets and noise models should be presented in order to improve the generalizability and robustness of the machine learning models and to enable a better and more insightful evaluation. In particular, the datasets should represent both sexes equally and include subjects of different ages, physiques, and with different heart diseases, as these are all limitations of the currently available datasets. Additionally, the noise model should focus on modeling artifacts that arise from real-world situations, such as from different activities and environments, but also other perturbations arising from the sensor electronics, for example. Finally, the performance of the denoising can be further improved by modifying the pre-processing or the architecture of the AEs or neural networks based on the findings presented in this work.

REFERENCES

[1] Aditya Sundar and Vivek Pahwa. 2017. Evaluating the Performance of State of the Art Algorithms for Enhancement of Seismocardiogram Signals. In *Proceedings of the First International Conference on Intelligent Computing and Communication (Advances in Intelligent Systems and Computing)*, Jyotsna Kumar Mandal, Suresh Chandra Satapathy, Manas Kumar Sanyal, and Vikrant Bhateja (Eds.). Springer, Singapore, 37–45. https://doi.org/10.1007/978-981-10-2035-3_5

[2] Manjula B M and Chirag Sharma. 2018. *Ballistocardiogram Signal Denoising Using Independent Component Analysis*. 259–267. https://doi.org/10.1007/978-981-10-4762-6_24

[3] Hsin-Tien Chiang, Yi-Yen Hsieh, Szu-Wei Fu, Kuo-Hsuan Hung, Yu Tsao, and Shao-Yi Chien. 2019. Noise Reduction in ECG Signals Using Fully Convolutional Denoising Autoencoders. *IEEE Access* 7 (2019), 60806–60813. https://doi.org/10.1109/ACCESS.2019.2912036

[4] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. 2016. Fast and Accurate Deep Network Learning by Exponential Linear Units (ELUs). arXiv:1511.07289 [cs]

[5] M. Di Rienzo, E. Vaini, P. Castiglioni, G. Merati, P. Meriggi, G. Parati, A. Faini, and F. Rizzo. 2013. Wearable Seismocardiography: Towards a Beat-by-Beat Assessment of Cardiac Mechanics in Ambulant Subjects. *Autonomic Neuroscience* 178, 1 (Nov. 2013), 50–59. https://doi.org/10.1016/j.autneu.2013.04.005

[6] Miguel A García-González, Ariadna Argelagós-Palau, Mireya Fernández-Chimeno, and Juan Ramos-Castro. 2013. A Comparison of Heartbeat Detectors for the Seismocardiogram. In *Computing in Cardiology 2013*. 461–464.

[7] Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. 2000. PhysioBank, PhysioToolkit, and PhysioNet: Components of a New Research Resource for Complex Physiologic Signals. *Circulation* 101, 23 (June 2000). https://doi.org/10.1161/01.CIR.101.23.e215

[8] Marian Haescher, Florian Höpfner, Wencke Chodan, Dimitri Kraft, Mario Aehnelt, and Bodo Urban. 2020. Transforming Seismocardiograms Into Electrocardiograms by Applying Convolutional Autoencoders. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. 4122–4126. https://doi.org/10.1109/ICASSP40776.2020.9053130

[9] Abdul Q. Javaid, Hazar Ashouri, Alexis Dorier, Mozziyar Etemadi, J. Alex Heller, Shuvo Roy, and Omer T. Inan. 2017. Quantifying and Reducing Motion Artifacts

in Wearable Seismocardiogram Measurements During Walking to Assess Left Ventricular Health. *IEEE Transactions on Biomedical Engineering* 64, 6 (June 2017), 1277–1286. https://doi.org/10.1109/TBME.2016.2600945

[10] Matti Kaisti, Mojtaba Jafari Tadi, Olli Lahdenoja, Tero Hurnanen, Antti Saraste, Mikko Pänkäälä, and Tero Koivisto. 2019. Stand-Alone Heartbeat Detection in Multidimensional Mechanocardiograms. *IEEE Sensors Journal* 19, 1 (Jan. 2019), 234–242. https://doi.org/10.1109/JSEN.2018.2874706

[11] Puneet Kumar Jain and Anil Kumar Tiwari. 2016. A Novel Method for Suppression of Motion Artifacts from the Seismocardiogram Signal. In *2016 IEEE International Conference on Digital Signal Processing (DSP)*. 6–10. https://doi.org/10.1109/ICDSP.2016.7868504

[12] David J. Lin, Jacob P. Kimball, Jonathan Zia, Venu G. Ganti, and Omer T. Inan. 2022. Reducing the Impact of External Vibrations on Fiducial Point Detection in Seismocardiogram Signals. *IEEE Transactions on Biomedical Engineering* 69, 1 (2022), 176–185. https://doi.org/10.1109/TBME.2021.3090376

[13] Loc Luu and Anh Dinh. 2018. Artifact Noise Removal Techniques on Seismocardiogram Using Two Tri-Axial Accelerometers. *Sensors (Basel, Switzerland)* 18 (04 2018). https://doi.org/10.3390/s18041067

[14] Lis Neubeck, Genevieve Coorey, David Peiris, John Mulley, Emma Heeley, Fred Hersch, and Julie Redfern. 2016. Development of an Integrated E-Health Tool for People with, or at High Risk of, Cardiovascular Disease: The Consumer Navigation of Electronic Cardiovascular Tools (CONNECT) Web Application. *International Journal of Medical Informatics* 96 (Dec. 2016), 24–37. https://doi.org/10.1016/j.ijmedinf.2016.01.009

[15] World Health Organization. 2023. *World Health Statistics 2023: Monitoring Health for the SDGs, Sustainable Development Goals*. World Health Organization.

[16] Keya Pandia, Sourabh Ravindran, Randy Cole, Gregory Kovacs, and Laurent Giovangrandi. 2010. Motion artifact cancellation to obtain heart sounds from a single chest-worn accelerometer. *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 590–593. https://doi.org/10.1109/ICASSP.2010.5495553

[17] Deepak Rai, Hiren Kumar Thakkar, Shyam Singh Rajput, Jose Santamaria, Chintan Bhatt, and Francisco Roca. 2021. A Comprehensive Review on Seismocardiogram: Current Advancements on Acquisition, Annotation, and Applications. *Mathematics* 9, 18 (Jan. 2021), 2243. https://doi.org/10.3390/math9182243

[18] D.M. Salerno and J.M. Zanetti. 1990. Seismocardiography : A New Technique for Recording Cardiac Vibrations. Concept, Method, and Initial Observations. *Seismocardiography : a new technique for recording cardiac vibrations. Concept, method, and initial observations* 9, 2 (1990), 111–118.

[19] Amirtahà Taebi, Brian E. Solar, Andrew J. Bomar, Richard H. Sandler, and Hansen A. Mansy. 2019. Recent Advances in Seismocardiography. *Vibration* 2, 1 (March 2019), 64–86. https://doi.org/10.3390/vibration2010005

[20] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. 2008. Extracting and Composing Robust Features with Denoising Autoencoders. In *Proceedings of the 25th International Conference on Machine Learning (ICML '08)*. Association for Computing Machinery, New York, NY, USA, 1096–1103. https://doi.org/10.1145/1390156.1390294

[21] Chenxi Yang, Foli Fan, Nicole Aranoff, Philip Green, Yuwen Li, Chengyu Liu, and Negar Tavassolian. 2021. An Open-Access Database for the Evaluation of Cardio-Mechanical Signals From Patients With Valvular Heart Diseases. *Frontiers in Physiology* 12 (Sept. 2021), 750221. https://doi.org/10.3389/fphys.2021.750221

[22] Chenxi Yang and Negar Tavassolian. 2016. Motion Noise Cancellation in Seismocardiogram of Ambulant Subjects with Dual Sensors. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 5881–5884. https://doi.org/10.1109/EMBC.2016.7592066

[23] J.M. Zanetti and D.M. Salerno. 1991. Seismocardiography: A Technique for Recording Precordial Acceleration. In *[1991] Computer-Based Medical Systems@m_Proceedings of the Fourth Annual IEEE Symposium*. 4–9. https://doi.org/10.1109/CBMS.1991.128936

APPENDIX

	-1	0	2	5	7	all
Train/Test Data	4.60	4.26	3.26	1.53	0.07	2.74
Results	4.99	4.57	3.29	1.33	-0.22	2.79
Train/Test Data + Results	4.78	4.40	3.21	1.28	-0.30	2.67
No Usage	**5.54**	**5.23**	**4.31**	**2.70**	**1.41**	**3.84**

Table 1: SNR improvements for usage of butterworth filters, see figure 2

	-1	0	2	5	7	all
Train/Test Data	0.09	0.08	0.07	0.06	0.06	0.07
Results	0.09	0.08	0.08	0.07	0.07	0.08
Train/Test Data + Results	0.09	0.09	0.08	0.07	0.07	0.08
No Usage	**0.08**	**0.07**	**0.06**	**0.05**	**0.05**	**0.06**

Table 2: RMSE for usage of butterworth filters, see figure 2

	-1	0	2	5	7	all
Train/Test Data	0.80	0.82	0.86	0.90	0.91	0.86
Results	**0.81**	**0.83**	**0.87**	**0.91**	**0.92**	**0.87**
Train/Test Data + Results	**0.81**	**0.83**	**0.87**	**0.91**	**0.92**	**0.87**
No Usage	0.80	0.82	0.86	0.90	0.92	0.86

Table 3: pcorr for usage of butterworth filters, see figure 2

	-1	0	2	5	7	all
None	3.14	2.50	0.71	-1.92	-3.77	0.13
WGN1	-1.49	-2.21	-4.04	-6.76	-8.67	-4.63
WGN2	2.60	2.10	0.77	-1.41	-3.11	0.19
WGN3	**5.54**	**5.23**	**4.31**	**2.70**	**1.41**	**3.84**
WGN4	4.62	4.26	3.25	1.51	0.13	2.75

Table 4: SNR improvements for different noise distributions, see figure 3

	-1	0	2	5	7	all
None	0.10	0.10	0.10	0.09	0.09	0.10
WGN1	0.15	0.15	0.14	0.14	0.14	0.14
WGN2	0.09	0.09	0.08	0.07	0.07	0.08
WGN3	**0.08**	**0.07**	**0.06**	**0.05**	**0.05**	**0.06**
WGN4	**0.08**	0.08	0.07	0.06	**0.05**	0.07

Table 5: RMSE for different noise distributions, see figure 3

	-1	0	2	5	7	all
None	0.65	0.67	0.74	0.81	0.85	0.74
WGN1	0.71	0.72	0.75	0.78	0.79	0.75
WGN2	0.78	0.80	0.84	0.88	0.89	0.84
WGN3	**0.80**	**0.82**	**0.86**	**0.90**	**0.92**	**0.86**
WGN4	0.79	**0.82**	**0.86**	**0.90**	**0.92**	**0.86**

Table 6: pcorr for different noise distributions, see figure 3

	-1	0	2	5	7	all
AE1	-2.2	-2.92	-4.74	-7.47	-9.36	-5.34
AE2	**5.54**	**5.23**	**4.31**	**2.7**	1.41	**3.84**
WD	3.43	3.24	2.72	1.94	**1.48**	2.56

Table 7: SNR improvements for different denoising methods, see figure 4

	-1	0	2	5	7	all
AE1	0.16	0.16	0.15	0.15	0.15	0.15
AE2	**0.08**	**0.07**	**0.06**	**0.05**	**0.05**	**0.06**
WD	0.10	0.10	0.08	0.06	**0.05**	0.08

Table 8: rmse for different denoising methods, see figure 4

	-1	0	2	5	7	all
AE1	0.76	0.78	0.82	0.86	0.88	0.82
AE2	**0.80**	**0.82**	**0.86**	**0.90**	**0.92**	**0.86**
WD	0.69	0.72	0.80	0.88	0.91	0.80

Table 9: pcorr for different denoising methods, see figure 4

	-1	0	2	5	7	all
D1 Dataset	**4.97**	**4.67**	**3.48**	**1.50**	**-0.16**	**2.88**
D1 Dataset [a]	4.71	4.38	3.24	1.24	-0.27	2.64

Table 10: SNR improvements for different datasets, see fig 5

	-1	0	2	5	7	all
D1 Dataset	**0.09**	**0.09**	**0.08**	**0.07**	**0.07**	**0.08**
D1 Dataset[a]	0.10	**0.09**	**0.08**	**0.07**	**0.07**	**0.08**

Table 11: rmse for different dataset, see figure 5

	-1	0	2	5	7	all
D1 Dataset	**0.83**	**0.85**	**0.89**	**0.92**	**0.94**	**0.89**
D1 Dataset [a]	0.79	0.81	0.85	0.89	0.91	0.85

Table 12: pcorr for different datasets, see figure 5

[a] Without Subject 1 and 2

Received 1 June 2023; revised x July 2023; accepted x July 2023