

Feasibility of ECG Reconstruction from Minimal Lead Sets Using Convolutional Neural Networks

Maksymilian Matyschik, Henry Mauranen, Pietro Bonizzi, Joël Karel

Maastricht University, Maastricht, the Netherlands

Abstract

Electrocardiography is a commonly applied method of measuring the electrical activity of the heart. The standard 12-lead electrocardiogram (ECG) provides sufficient information to allow various heart conditions to be diagnosed. Despite its relative ease of use, the standard ECG procedure could benefit from a reduction of leads which may allow for continuous monitoring, for example via a wearable device. In this study, we first investigated the use of variational autoencoders (VAEs) to assess what the most representative leads of the standard 12-lead system are. As the VAE learns to compress the ECG data, it focuses on the parts of the input that is important for the reconstruction. This information is then used to assess which leads are the most useful for a reconstruction task in general. Precordial leads V_2 , V_3 and V_4 are shown to contain the most information in the 12-lead ECG data. We then investigated the use of a convolutional neural network (CNN) architecture capable of learning patient-specific models to accurately impute 11 missing ECG signals from a single available lead. Our design is unconventional in that it keeps a two-dimensional structure throughout the fully connected layers. We show that this design outperforms the traditional one-dimensional structure and that these architectures can be affected by the presence of symptoms in recorded heart signals.

1. Introduction

Reconstruction of 12-lead ECGs from a minimal number of leads can be relevant for wearables and continuous monitoring. Despite the procedure being widely used and well standardized, it is challenging to perform continuously during normal day-to-day life of a subject. [1] investigated lead reduction in 64-lead body surface potential maps (BSPMs). Other ECG procedures have been shown to contain redundant information [2], which shows promising results for applying neural networks in compressing and reconstructing ECG information. [2] implemented neural networks in the context of 65-lead BSPMs

and significantly reduced the lead count required for reliable reconstructions.

This raises the question if similar approaches could be applied to the traditional 12-lead ECG to achieve more portable monitoring system. In [2], a heuristic method based on VAE is used for lead selection and assessment of lead significance for the reconstruction process. In this study VAEs were analyzed directly to understand which leads had highest contribution to reconstruction quality. Further, we assessed the performance of two convolutional neural network (CNN) [3] architectures in the context of patient-specific 12-lead ECG reconstructions from a single lead. This analysis was performed on a dataset containing healthy and patient recordings. We investigated whether single-lead reconstructions are a viable possibility, which architecture is superior and compared reconstruction quality across the various subsets of recordings in the dataset.

2. Methodology

2.1. Data

Two datasets were used during this study. For VAE experiments, PTB Diagnostic database was used [4]. The dataset contains both healthy subjects and patient records, sampled at 1000 Hz. A high-pass Butterworth filtering was applied with cut-off frequency of 0.5 Hz to eliminate baseline wander. The end effects of the filtering were then eliminated by omitting short segments from the beginning and end of each recording.

The T-Wave Alternans Challenge database (TWA)[5] was used for CNN reconstruction. This dataset contains 100 records of healthy subjects and patients (sampled at 500 Hz) showing symptoms of myocardial infarction and other conditions, as well as two synthetic subsets featuring healthy signals and modified data with artificially added T-wave alternans. The recordings were preprocessed with a second-order pass-band Butterworth filter with critical frequencies 0.5 Hz and 50 Hz to remove breathing frequencies and other sources of noise, such as powerline interference. Each recording was divided into windows of 512 timesteps with a stride of 100 to allow for easier process-

ing with the CNN.

Finally, the data in both cases were scaled to $[-1, 1]$ over all channels, and split into training and testing sets; the latter were taken from the end of the recordings for patient-specific models. VAE applied a 75/25 split between training and testing data while CNN used approximately 10% for testing and divided the remaining data into training and validation sets using an 80/20 ratio. Windows that overlap the test and train sets were discarded.

2.2. Network Architectures

2.2.1. Investigation of Lead Significance

Convolutional VAEs (ConVAEs) based on the structure in [2] were used to investigate the lead significance and amount of dimensionality reduction on the input data. For this purpose, the investigated hyperparameters were the repetition of the convolutional blocks, dense blocks and the size of the latent layer. As the purpose of the experiments was to analyze a VAE and what it learns, the structure was kept simple, LeNet based, instead of using architectures like XceptionNet [6] that might offer better reconstruction performance.

The input data was windowed for these experiments based on the parameters found in [2]. This ensures manageable training times without compromising the reconstruction quality. Windows of 2048 time steps (2 seconds) were used, with 512 step overlap. These values allow repetition of the convolution blocks without having to change the window size. The small overlap was used to control the size of the data and overfitting.

Gradient-weighted class activation mapping (Grad-CAM) [7] is typically used to examine the sections of input that result in a specific class. By examining small gradients, rather than the usual positive ones, we can highlight the areas responsible for the most correct regression value. This way it becomes possible to visualize the areas of the input that are responsible for a good reconstruction and visualize both the most important leads and time segments of the ECG. For the latter, it can be expected to be QRS-complex and T-wave.

2.2.2. Reconstructions of 12-Lead ECGs from One Lead

We devised two CNN architectures to reconstruct missing ECG data. This choice was made based on our previous work [2]. Existing architectures such as AlexNet [8] were used as inspiration for our design due to their performance in image-related tasks. We opted to apply one-dimensional convolutions since we are dealing with time-

series data. Conventional CNNs flatten the output of the convolutional block and pass it through a number of fully connected layers. This structure can be found in the architecture we dubbed "CF". Furthermore, we added a second architecture which does not flatten the intermediate output but instead reshapes it into a two-dimensional matrix that allows its vectors to be processed by the fully connected layers one by one. We named this architecture "CR". Each of these models was additionally tested in a variation containing dropout layers with a dropout rate of 10% to add a degree of regularization [9]. These models will be referred to as CR-D and CF-D.

Our CNNs were compared against a linear baseline model. This least squares approach (LS) [1] applies a transfer matrix \mathbf{T} to the recorded input signals \mathbf{X}_{in} to reconstruct the missing signals $\hat{\mathbf{X}}_{out}$. \mathbf{T} is found using all 12 channels in the training phase by applying $\mathbf{T} = \mathbf{X}_{out}\mathbf{X}_{in}^\dagger$. The Moore-Penrose pseudoinverse \mathbf{X}_{in}^\dagger is used to compute the least squares solution. Then, the missing signals can be obtained as follows: $\hat{\mathbf{X}}_{out} = \mathbf{T}\mathbf{X}_{in}$. One model was trained for each patient and input lead. 45 test windows were extracted from each patient's recording. We used Mean Squared Error (MSE) to assess the models' performance.

3. Results

3.1. VAE Experiment Results

ConVAE is able to reconstruct the 12-lead ECG reasonably well. The most prominent features are well reconstructed in the QRS-complex and T-wave, with some issues in the leads where these features are less prominent. Some of the R-peak heights are not perfectly reached, but the shape and structure is correctly present. As the network achieves this quality of reconstructions, it can be expected to have learned the shapes and structures present in the 12-lead ECG recording to a certain extent.

Grad-CAM is applied to all windows from all healthy subjects and the results are summed and normalized to obtain relative significance of each lead, in Figure 1. The scaling is based on percentage of the total contribution. Here we observe that leads V_2 , V_3 and V_4 are considered to be very important by the network, while it has decided to include almost none of the information from leads such as aVR and III . Similar figures are obtained for different heart conditions as well, which also emphasize the precordial leads. In some of the patient results, II came out with slightly higher significance than in healthy cases.

3.2. CNN

Our results indicate that CR outperforms the other models to a significant degree (Figure 2). The addition of dropout layers did not improve reconstruction quality.

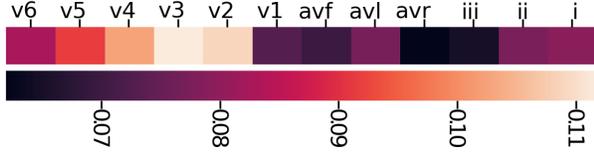


Figure 1. Relative significance of each lead for the reconstruction in ConVAE. The legend on lower row displays fraction of total (1.0) importance for the reconstruction.

Both variations of CF were outperformed by CR. The linear baseline model yielded the highest MSE. The best reconstruction of a test window outputted by CR is shown in Figure 3. In this setup, the network was trained on lead index 8 which corresponds to precordial lead V_3 (highlighted by green background). The patient’s data was taken from the PTB database. The blue lines represent the recorded signals; predicted data is shown in red.

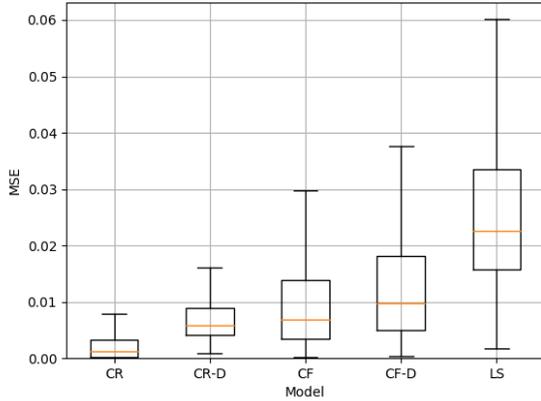


Figure 2. Comparison of All Models, Healthy and Patient Data, MSE

Figure 4 displays a comparison of a healthy control group (PTBc) and a set of patients (PTB). Outliers appear as red dots. The healthy control groups yielded lower MSE for both CR and CF. However, this difference is more visible in CR’s metrics, especially when outliers are taken into account.

3.3. Discussion

The ConVAE takes in the 12-lead ECG recordings, compresses the information present to the latent layer, and reconstructs the original input. In order to do this, the network needs to learn compact representations of the input data, and must be able to focus on the more important regions of the input. As the experiments display that the ConVAE is able to perform the reconstruction to a reasonable degree, important regions for the reconstruction are also the information rich regions of the ECG process. In

terms of leads, the most significant one can be seen to be V_3 . While this is unsurprising from intuitive perspective, the lead being located on top of the heart, the network has originally not received any geometric information about the lead positions.

This experiment has shown the significances of the individual leads for achieving overall good reconstructions. These results slightly change for different heart conditions, but do not allow to understand if a particular lead is significant for only some heart conditions. A classifier network could be used to unveil such information. Further research could also be done towards the latent space. For this ConVAE structure, the latent space did not provide fruitful results, but more complex networks might yield useful and compact descriptions of the ECG process as whole.

If no major anomalies occur in the test data, CR can be expected to produce reliable reconstructions of 12-lead ECGs from a single channel. Care must be taken to select a viable lead for each patient, however. We observed a trend towards precordial leads in this context.

Patient data is more difficult to reconstruct; this difference is more visible in CR’s metrics. It could be argued that this issue can be mitigated as long as the symptoms are somewhat consistent and well-represented in the training data. By contrast, novel deviations will yield a significantly greater error due to the fact that the personalized model was not able to learn the new pattern. Despite this limitation, the network may still be capable of producing signal shapes that clearly indicate an anomaly; this information could then be used by a classifier or a trained physician to identify the presence of a heart condition.

The unconventional two-dimensional structure in the fully connected block of CR led to better results than the ones produced by the conventional architecture used in CF. We speculate this may be related to the inherent regularization attributed to the reduced parameter count and the weight sharing which is a result of multiple vectors being processed by the fully connected layers with identical parameters.

4. Conclusion

The correct lead selection and placement is an important step towards portable ECG devices. In this study, it was shown that the precordial leads, in particular V_2 , V_3 and V_4 , are promising candidates for this in a neural network based approaches. While this follows the intuitive placement of the leads, this study shows that it is supported by the measured data. We conclude that CR is a viable architecture for personalized 12-lead ECG reconstruction tasks with one available input lead. This finding indicates that a substantial lead reduction in ECG procedures can be achieved, producing results that closely mirror the true heart signals

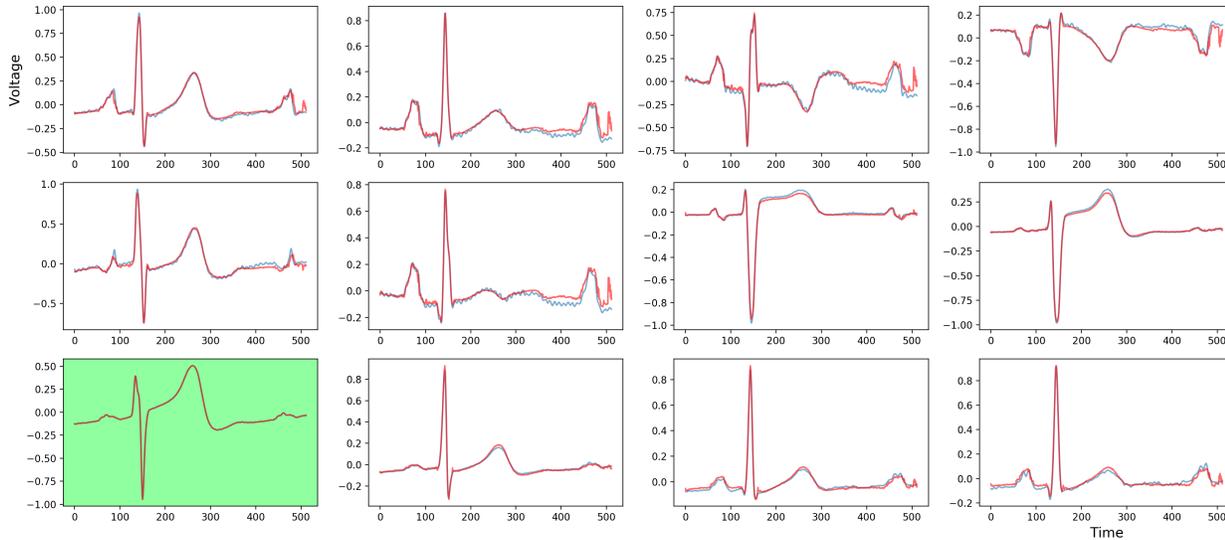


Figure 3. Best Reconstructions by CR, Patient twa11; Red Lines Represent CNN Output, Input Lead V_3 is Highlighted in Green

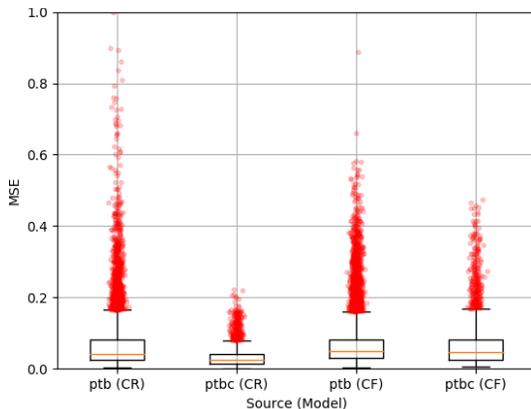


Figure 4. Comparison of Healthy and Patient Data Reconstructions, MSE

as long as no considerable abnormalities occur. Future research may investigate the feasibility of this approach in a clinical context, with particular focus on the potential difficulties resulting from the patient's movements during continuous monitoring via a wearable device.

References

[1] Guillem M, Bollmann A, Climent AM, Husser D, Millet-Roig J, Castells Ramon F. How many leads are necessary for a reliable reconstruction of surface potentials during atrial fibrillation? *IEEE transactions on information technology in biomedicine a publication of the IEEE Engineering in Medicine and Biology Society* 06 2009;13:330–40.

[2] Bizcaino D, Bujnarowski K, Matyschik M, Mauranen H, Zhao R, Bonizzi P, Karel J. Neural network-based matrix

completion for minimal configuration of bodysurface potential mapping. In *Computing in Cardiology*. 2019; .

[3] Lecun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 12 1998;86:2278 – 2324.

[4] Bousseljot R, Kreiseler D, Schnabel A. Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet. 1995; .

[5] Moody G. The physionet / computers in cardiology challenge 2008: T-wave alternans, 2008.

[6] Chollet F. Xception: Deep learning with depthwise separable convolutions. *CoRR* 2016;abs/1610.02357. URL <http://arxiv.org/abs/1610.02357>.

[7] Selvaraju RR, Cogswell M, Das A, Vedantam R, Parikh D, Batra D. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *2017 IEEE International Conference on Computer Vision (ICCV)*. ISSN 2380-7504, Oct 2017; 618–626.

[8] Krizhevsky A, Sutskever I, Hinton G. Imagenet classification with deep convolutional neural networks. *Neural Information Processing Systems* 01 2012;25.

[9] Srivastava N, Hinton G, Krizhevsky A, Sutskever I, Salakhutdinov R. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research* 06 2014;15:1929–1958.

Address for correspondence:

Pietro Bonizzi
 P.O. Box 616, 6200 MD Maastricht, The Netherlands
 pietro.bonizzi@maastrichtuniversity.nl