

Deep Convolutional Encoder-Decoder Framework for Fetal ECG Signal Denoising

Eleni Fotiadou¹, Tomasz Konopczyński², Jürgen Hesser², Rik Vullings¹

¹ Department of Electrical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

² Department of Radiation Oncology, Medical Faculty Mannheim, Central Institute of Scientific Computing (IWR), and Central Institute for Computer Engineering (ZITI), Heidelberg University, Germany

Abstract

Non-invasive fetal electrocardiography has the potential of providing vital information for evaluating the health status of the fetus. However, the low signal-to-noise ratio of the fetal electrocardiogram (ECG) impedes the applicability of the method in clinical practice. Residual noise in the fetal ECG, after the maternal ECG is suppressed, is often non-stationary, complex and has spectral overlap with the fetal ECG. We present a deep fully convolutional encoder-decoder framework, for removing the residual noise from single-channel fetal ECG. The method was tested in a broad simulated fetal ECG dataset with varying amount of noise. The results demonstrate that after the denoising there was an average increase in the correlation coefficient between the corrupted signals and the original ones from 0.6 to 0.8. Moreover, the suggested framework successfully handled different levels of noises in a single model. The network was further tested on real signals showing substantial noise removal performance, thus providing a promising approach for fetal ECG signal denoising. The presented method is able to significantly improve the quality of the extracted fetal ECG signals, having the advantage of preserving beat-to-beat morphological variations.

1. Introduction

During pregnancy and labor, non-invasive monitoring of the fetal heart is of great clinical interest for assessing the condition of the fetus. Non-invasive fetal electrocardiography, performed by abdominal electrodes, is a technology that has significantly evolved the last years [1]. However, it still needs a great deal of improvements to become a valid alternative to cardiotocography, the technique commonly used in

clinical practice. Nevertheless, fetal electrocardiography holds the promise of providing additional information in terms of the ECG morphology. The reason for this required improvement is that the abdominal signals are severely contaminated by the maternal ECG, muscle noise, baseline wander, powerline interference and other interferences and noises. The majority of these noises overlap in frequency and time with the fetal ECG, rendering the extraction of a clean fetal ECG a challenging signal processing task.

According to the review of Sameni *et al.* [1] the main methods for fetal ECG extraction include adaptive filtering and linear and non-linear decomposition. In the extracted fetal ECG signals, usually it is possible to detect the QRS complex due to the high amplitude of the fetal R-peak. However, the low signal-to-noise-ratio (SNR) of the fetal ECG impedes further morphological analysis and post-processing to remove residual noise becomes inevitable. The running average of several heartbeats is widely used to improve the SNR of the signal at the expense of losing individual variations in pulse shape. In a previous work, the authors have used an augmented time-sequenced adaptive filter as a post-processing quality enhancement step [2]. However, the fetal pulse locations are needed to synchronise the filter and additionally the filter is unable to track rapid changes in the signal morphology, for instance in the presence of ectopic beats.

To overcome the aforementioned problems, we propose the use of a deep learning approach for single-channel fetal ECG signal denoising. Stacked denoising autoencoders [3] and recurrent neural networks [4] have been already suggested for denoising adult ECG. Instead, inspired by the work of Mao *et al.* [5] on image restoration, we propose the employment of a deep convolutional encoder-decoder network with symmetric skip connections.

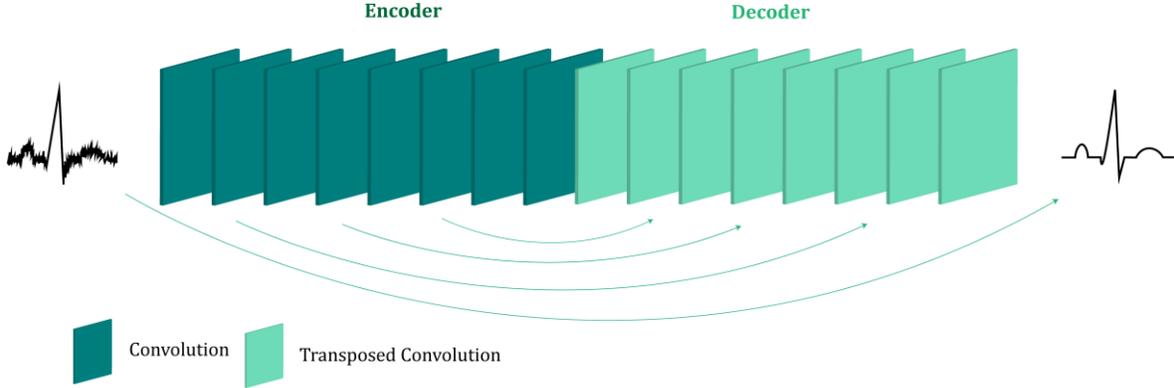


Figure 1. The architecture of the proposed fetal ECG denoising network.

2. Methods

2.1. Network Architecture

The proposed network consists of a chain of convolutional and symmetric transposed convolutional layers, as shown in Figure 1. The encoder part of the network acts as a feature extractor, which captures the primary fetal ECG structure, while eliminating the signal corruptions layer by layer. During this process the signal details might be lost. Then, the decoder part is used to compensate for the missing details. The decoder aims to deliver the clean version of the input signal as output. Moreover, skip connections are added between every two corresponding convolutional and transposed convolutional layers. Through the skip connections, the convolutional feature maps are added to the symmetric transposed convolutional ones and then pass to a Leaky Relu layer. The role of the skip connections is to transfer signal details to the decoder, aiding in recovering the clean signal. Moreover, they help in tackling the optimization difficulty caused by gradient vanishing in deep architectures.

Our network comprises of 8 convolutional and 8 transposed convolutional layers. All the convolutional layers, apart from the first one, include subsampling operations. For denoising tasks down-sampling of the signal is not generally preferred for the sake of preserving the signal details. However, by employing down-sampling operations the effective receptive field of the network is significantly increased. This is particularly important in the case of fetal ECG as it usually exhibits high levels of noise and thus it requires a large effective patch size to capture more information for efficient denoising. As a kernel size we use 15 which in the case of signals sampled at 500Hz corresponds to 30ms. The above parameter choices lead to an effective receptive field of approximately 3.6s that corresponds to 5-10

heartbeats. The number of feature maps in the convolutional layers is [64, 128, 128, 256, 256, 512, 512, 1024]. Consequently, the dimension of the bottleneck feature vector is 1024x15. The transposed convolutional layers use mirrored numbered of filters.

As a loss function the normalized mean squared error (NMSE) is adopted:

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \frac{(X^{(i)} - \hat{X}^{(i)})^2}{\bar{X}^{(i)2}} \quad (1)$$

where N is the number of the training data in a batch, $\hat{X}^{(i)}$ is the i-th denoised fetal ECG, $X^{(i)}$ is the clean version of $\hat{X}^{(i)}$, used as ground truth, and $\bar{X}^{(i)}$ the mean signal amplitude. For optimizing the network, the Adam algorithm is selected.

2.2. Performance metric

The performance of the method is assessed based on the correlation coefficient between the network input and output:

$$R_{ij} = \frac{c_{ij}}{\sqrt{c_{ii} * c_{jj}}} \quad (2)$$

where C is the covariance matrix. The values of R are between -1 and 1 inclusive.

3. Data

For training of the proposed network as well as for testing we use a simulated dataset. Moreover, to show the network potential in real cases we include also results of real data.

The simulated dataset is composed of two parts. The first part of the simulated dataset consists of the Fetal ECG Synthetic database of Physionet (FECGSYNDB) [6], [7]. The FECGSYNDB includes 1750 synthetic 34-channel abdominal signals, sampled at 250 Hz and each 5 minutes long. Different physiological events are simulated, such as fetal movements, heart rate accelerations/ decelerations, ectopic beats etc. To increase

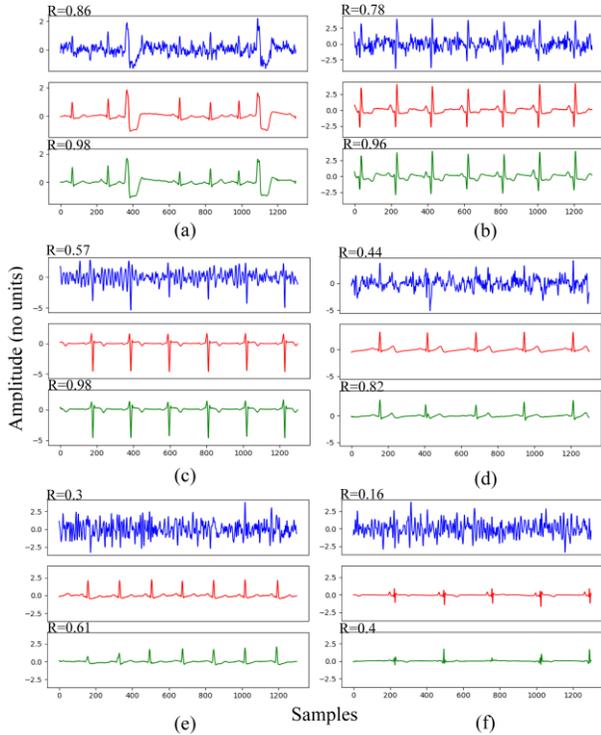


Figure 2. Denoising results for simulated fetal ECG signals in the test dataset. For each subplot a-f: the noisy signal is presented in the first row (blue), followed by the ground truth in the second one (red) and the denoised signal in the last row (green). The correlation coefficients, R , between the noisy/denoised signals and the ground truth are also provided.

the diversity of ECG patterns in our dataset and reduce the risk of overfitting we use a second set of data comprising of adult ECG and real noise. The adult ECGs are collected from 3 different databases of Physionet, the PTB Diagnostic ECG Database [8], the St.-Petersburg Institute of Cardiological Technics 12-lead Arrhythmia Database [7] and the QT database [9]. The signals are denoised and resampled at half frequency to resemble fetal ECG. Afterwards, noise is added to the signals using a set of abdominal measurements of an ongoing study of which the study protocol is described in [10]. In some of these measurements, the fetal ECG was impossible to be detected either because of the shielding of the fetus by the vernix caseosa or because some electrodes were far from the fetal heart. These cases, after suppression of maternal ECG and powerline interference, were considered as pure noise and added to the adult ECG for fetal ECG simulation. The SNR of the simulated dataset ranges from -15 to 15dB. Finally, the signals are separated in training and test data, ensuring that there is no overlapping between them. The real fetal ECG was recorded on 6 channels by the Nemo Healthcare fetal monitor, operating at 500 Hz sampling rate, as described in [10].

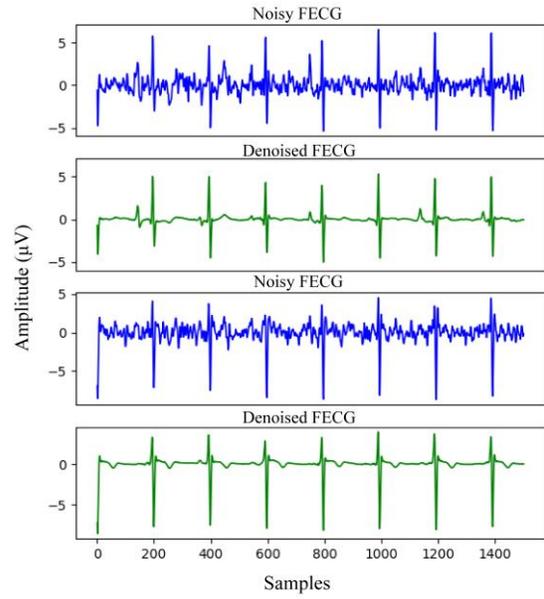


Figure 3. Denoising result in 2 channels of a real non-invasive fetal ECG recording.

The signals of all datasets were pre-processed before entering the network, either for training or for testing. The fetal ECG extraction is performed by using the open-source algorithm of Varanini [11] and all the signals are resampled to 500Hz to have a common reference. Then, they are divided in segments of 1920 samples and finally are normalized to have zero mean and unity standard deviation.

4. Results and Discussion

The proposed network is trained until convergence is reached and subsequently evaluated in the test dataset. The correlation coefficient measured between the corrupted and the clean signals increased to 0.8 when compared to 0.6 prior to denoising. Some representative denoising results in the test signals are shown in Figure 2. As can be seen in the figure, for cases of lower noise (Figure 2(a), (b) and (c)) the network not only suppresses the noise significantly but also reveals the ECG waves to an extent that the similarity with the ground truth is very high, resulting in a correlation coefficient close to 1. When the amount of noise is higher (Figure 2(d) and (e)) the QRS complexes are usually recovered but there might be some distortion in the signal. The smaller waves, like the P and T wave, are more difficult to be reconstructed reliably. Finally, when the SNR of the signal is really low (Figure 2(f)) even the QRS complexes appear difficult to recover. It should be pointed out that the network works surprisingly well in cases of arrhythmia as can be observed in Figure 1(a). This means that the network does not simply learn that successive complexes should be

rather similar but additionally learns how to preserve the variations in individual pulse shapes.

Figure 3 demonstrates a result when denoising two of the channels of a 6-channel real fetal ECG. The denoised signals are free from noise and there is some correspondence between the waves of the different channels. However, since there is no ground truth signal available to compare, there is no guarantee that the morphological features of the fetal ECG were correctly reconstructed. To confirm that the method is able to work also in real signals, validation with simultaneous scalp measurements is necessary. Even though the scalp ECG is a different ECG lead than the abdominal leads and they should not be identical, in terms of morphology, measured intervals between the denoised signals and the scalp ECG should have similar values.

Despite the limitation in our validation, this is the first study to demonstrate fetal ECG signal denoising with deep convolutional neural networks (CNNs). One of the main advantages of the method is that no prior processing of the signal is needed to extract the locations of the R-peaks as opposed to the widely adopted running averaging method. Moreover, individual variations among different ECG complexes are preserved. This is especially beneficial in arrhythmia cases. So far, arrhythmia is assessed through echocardiography because the averaging performed to enhance the quality of the fetal ECG results in loss of relative information. The results of this work show that the fetal ECG can become a tool to study arrhythmia cases. In addition, the quality of the signals is improved to the extent that can allow for measuring the exact timing of morphological features by the clinicians, like the PR and QT interval.

There is certainly room for improvement of our method. In cases of very low SNR the results might not be reliable since the signal is distorted and some waves are not reconstructed accurately or appear to have opposite polarity. It might be that the receptive field of our network is not large enough to capture enough information for denoising such highly corrupted signals. However, increasing the receptive field is computationally very demanding and thus there is a trade-off between performance and efficiency. Moreover, we currently use merely a single channel for denoising. Combining information from multiple channels might improve the performance of the network as it will allow correlations to extend also in space rather than only in time. Finally, including more variations in the ECG morphology of the dataset can be beneficial to make the network more generalizable.

6. Conclusion

The proposed deep encoder-decoder framework is able to achieve a substantial quality improvement of noisy fetal ECG signals. The principal advantage of the method

is that individual variations among different pulses can be preserved and that, opposed to most other fetal ECG denoising methods, the method does not require knowledge on R-peak locations.

References

- [1] R. Sameni and G. D. Clifford, "A Review of Fetal ECG Signal Processing: Issues and Promising Directions," *Open Pacing Electrophysiol Ther J.*, p. 4–20, 2010.
- [2] E. Fotiadou, J. O. E. H. van Laar, S. G. Oei and R. Vullings, "Enhancement of low-quality fetal electrocardiogram based on time-sequenced adaptive filtering," *Med. Biol. Eng. Comput.*, vol. 56, no. 12, p. 2313–2323, 2018.
- [3] P. Xiong, H. Wang, M. Liu, S. Zhou, Z. Hou and X. Liu, "ECG signal enhancement based on improved denoising auto-encoder," *Eng. Appl. Artif. Intell.*, vol. 52, pp. 194–202, 2016.
- [4] K. Antczak, "Deep Recurrent Neural Networks for ECG Signal Denoising," arXiv:1807.11551, 2018.
- [5] X. Mao, C. Shen and Y. Yang, "Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections," in *NIPS*, Barcelona, 2016.
- [6] F. Andreotti, J. Behar, S. Zaunseider, J. Oster and G. D. Clifford, "An open-source framework for stress-testing non-invasive foetal ECG extraction algorithms," *Physiol. Meas.*, vol. 37, no. 5, pp. 627–48, 2016.
- [7] A. L. Goldberger et al., "PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, no. 23, pp. e215–20, 2000.
- [8] R. Boussejot, D. Kreiseler and A. Schnabel, "Nutzung der EKG-Signaldatenbank CARDIODAT der PTB über das Internet," *Biomedizinische Technik*, vol. 40, no. s1, 1995.
- [9] P. Laguna, R. G. Mark, A. L. Goldberger and G. B. Moody, "A Database for Evaluation of Algorithms for Measurement of QT and Other Waveform Intervals in the ECG," in *Computers in Cardiology*.
- [10] K. M. J. Verdurmen, C. Lempersz, R. Vullings and C. Schroer, "Normal ranges for fetal electrocardiogram values for the healthy fetus of 18–24 weeks of gestation: a prospective cohort study," *BMC Pregnancy Childbirth*, vol. 16, no. 1, p. 227, 2016.
- [11] M. Varanini, G. Tartarisco, L. Billeci, A. Macerata, G. Pioggia and R. Balocchi, "A multi-step approach for non-invasive fetal ECG analysis," in *Computing in Cardiology 2013*, Zaragoza, 2013.

Address for correspondence:

Eleni Fotiadou
P.O. Box 513, 5600 MB Eindhoven, The Netherlands
e.fotiadou@tue.nl